

Brewer, Mike; Cominetti, Nye; Jenkins, Stephen P.

Working Paper

What Do We Know About Income and Earnings Volatility?

IZA Discussion Papers, No. 17808

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Brewer, Mike; Cominetti, Nye; Jenkins, Stephen P. (2025) : What Do We Know About Income and Earnings Volatility?, IZA Discussion Papers, No. 17808, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/316763>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 17808

**What Do We Know About Income and
Earnings Volatility?**

Mike Brewer
Nye Cominetti
Stephen Jenkins

MARCH 2025

DISCUSSION PAPER SERIES

IZA DP No. 17808

What Do We Know About Income and Earnings Volatility?

Stephen Jenkins

LSE and IZA

Mike Brewer

LSE

Nye Cominetti

Resolution Foundation

MARCH 2025

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

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

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

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

What Do We Know About Income and Earnings Volatility?*

We first review research about income and earnings volatility and second provide new UK evidence about the latter using high quality administrative record data. The USA stands out as a high volatility country relative to the UK and other high-income countries, but volatility levels have remained constant in these countries recently. Almost all research has considered volatility from an annual perspective whereas we provide new evidence about month-to-month earnings volatility. There is a distinct within-year seasonal pattern to volatility, and volatility is highest for the top and bottom tenths of earners. High earnings volatility among top earners and its seasonality reflect pay bonus patterns whereas, for low earners, the instability of hours including zero-hours contracts likely play important roles. Our findings have relevance to the design of cash transfer support in the UK because the monthly reference periods it uses do not align with many earners' pay periods.

JEL Classification: D31, I31, J31, J38

Keywords: income volatility, earnings volatility, PAYE data, administrative record data, survey data

Corresponding author:

Stephen P. Jenkins
Department of Social Policy
London School of Economics and Political Science
Houghton Street
London WC2A 2AE
United Kingdom
E-mail: s.jenkins@lse.ac.uk

* This article is based on Jenkins's Nancy and Richard Ruggles Memorial Lecture at the IARIW General Conference, London, August 2024. Section 5 draws on research, part-funded by the Joseph Rowntree Foundation, that piggybacks on the Work and Employment Dynamics (WED) project, funded by Administrative Data Research UK. Our calculations using HMRC's real-time information PAYE earnings data linked to data from the Annual Survey of Hours and Employment were undertaken in the Office for National Statistics Secure Research Service using data from the ONS and other owners and do not imply the endorsement of the ONS or other data owners. Special thanks to Van Phan and Felix Ritchie (WED) for their help creating datasets, for post-Lecture comments from Richard Blundell, Alex Bryson, Karen Dynan, John Forth, Peter Ganong, Fatih Guvenen, Alan Manning, Jonathan Morduch, Robert Moffitt, John Sabelhaus, and Carl Singleton, and comments from two anonymous referees and the handling editor (Robert Inklaar). This paper was written while Jenkins was a visiting fellow at the Department of Economics, European University Institute.

1. Introduction

Volatility refers to fluctuations over time of the incomes of individuals within some population. This paper first reviews what we know about the volatility of household income and individual employment earnings. Second, we provide new evidence about within-year earnings volatility in the UK and argue it indicates problems with the administration of the current social assistance benefit system.

We develop our arguments under five main headings, beginning with household income volatility. Most studies refer to data that are a decade old: there are few studies of income volatility trends within countries and few cross-national comparative studies using recent data. This is a problem because household income is widely accepted to be a better measure of household living standards than individual labour earnings and, with measures of household income, one can also examine the extent to which volatility in one member's labour earnings is offset (or reinforced) by other components in the household income package such as taxes and cash transfers, other labour earnings, and non-labour income. Evidence indicates that the USA is an outlier with high income volatility compared to other high-income countries following a marked rise in the 1970s and 1980s, though, as for other countries, more recent data indicate that income volatility remained broadly constant through to at least the middle of the 2010s.

Second, we consider volatility in employment earnings. There is more and more up-to-date information than there is for income. This is especially so for the USA but also for other countries including the UK. Relatedly, there is a growing use of administrative record data to study earnings volatility, whereas income volatility studies continue to be based on household panel surveys. Earnings volatility is relatively high in the USA compared to other countries but shows a flat trend over the last decade.

Third, not enough is known about the volatility of income or earnings within a year. Reflecting the nature of the longitudinal datasets most commonly available, almost all studies refer to annual volatility, by which we mean either changes in annual income between one year and the next (as in most US analyses), or changes in monthly income between pairs of months approximately a year apart (as in UK analyses). In both cases, within-year volatility is assumed away. An annual income definition smooths incomes over the year by construction. Comparing two monthly incomes a year apart misses income variation across the months in between. However, most people receive their incomes, whatever the source, at a monthly or higher frequency and so, naturally, there is substantial interest in how much within-year

volatility there is, who experiences it and, by implication, whether patterns reflect a potential policy problem.

Fourth and relatedly, we provide new evidence about the volatility of monthly employment earnings, month by month over the five years prior to Covid-19 onset. Our high quality and large-scale administrative record data for the UK reveal sizeable variations in earnings volatility within a year: there is a seasonal pattern to volatility that has received little attention. There are also large differences in monthly volatility levels across subgroups of earners with, for example, volatility greatest for those at the bottom and at the top of the earnings distribution but with different seasonal patterns. Pay bonus patterns likely account for high volatility among top earners, whereas for low earners the high prevalence of zero-hours contracts likely plays an important role, along with instability of work hours more generally.

Fifth, these findings are relevant to the design and administration of cash transfer (benefits) policy. We argue that, for many low-income UK families, high frequency volatility leads to problematic misalignment between the reference periods for income received from paid work and the reference period assumed by the cash benefit system.

Making this argument is not to say that all income volatility is welfare-reducing and of policy concern. Volatility is a signal of income risk, and individuals and families are typically assumed to be risk averse. As Shorrocks has put it, “[g]reater variability of incomes about the same average level is disliked by individuals who prefer a stable flow. So, to the extent that mobility leads to more pronounced fluctuations and more uncertainty, it is not regarded as socially desirable” (1978, pp. 392–3). However, the signal about risk is noisy. Some income changes may be anticipated (such as annual bonuses for high earners) or deliberately chosen, and families may insure against adverse changes by borrowing and saving, changing their labour supply, or be helped by the tax-benefit system.

Nonetheless, there is sufficient evidence that volatility causes problems for low-income families, the group most likely to be liquidity constrained, and of greatest social concern. Our new empirical evidence for the UK highlights substantial earnings volatility for this at-risk group which may be exacerbated by the cash transfer system intended to provide them with support. Volatility also leads to other problems which impact low-income households more than others, including the increased psychological, cognitive, and emotional costs of managing with unstable income, and reduced abilities to save or make longer-term financial investments in, e.g., pensions or housing. (See Angsten Clark and Otulana, 2024.)

We elaborate on the points under our five headings in the rest of the paper. For brevity, additional discussion, charts, and tables that supplement our analysis of UK monthly earnings volatility are placed in Appendices A–E (Online Supplementary Material). For additional analysis, including of week-to-week volatility patterns among weekly-paid workers, plus discussion of policy implications, see Brewer, Cominetti, and Jenkins (2025).

Throughout, we assume that the interest is in volatility per se, rather than the relationships between volatility and mobility, or how changes in volatility contribute to changes in overall inequality. We do not consider the more general concept of ‘economic insecurity’, about which see, e.g., Osberg (2015). Our discussion of empirical volatility estimates is restricted to high-income countries, primarily the USA and the UK. For low- and middle-income countries there is analysis of poverty and income dynamics but not of volatility, as far as we know. On the related topic of vulnerability to poverty in poor countries, see e.g. Merfeld and Morduch (2024) and references therein. Before reviewing existing findings about household income and individual earnings volatility, we explain how volatility is measured.

2. Volatility measures

Volatility measures can be classified into period-pair and multiperiod measures. Period-pair measures are based on changes in income between a pair of time periods for each person. A key analytical choice is how to summarize income ‘change’. The most common approach is to use the *arc percentage change*. For each person, $i = 1, \dots, N$, with current period t ’s income y_{it} and previous period income y_{it-1} , the arc percentage change, a_{it} , is

$$a_{it} = \frac{y_{it} - y_{it-1}}{(y_{it} + y_{it-1})/2} = \frac{y_{it} - y_{it-1}}{\bar{y}_i}$$

where \bar{y}_i is the longitudinal average of i ’s pair of incomes. The arc percentage change has nice properties. It is symmetric in income rises and falls, e.g., for an income rise from 50 to 100 and a fall from 100 to 50, the arc percentage changes are 67% and –67%, whereas the corresponding conventional percentage changes are 100% and –50%. Moreover, the arc percentage change allows us to calculate changes from zero to some positive value, and the reverse, generating upper and lower bounds of 200% and –200%. (By convention, if $y_{it-1} = y_{it} = 0$, $a_{it} = 0$.) We can therefore derive estimates not only for *earnings volatility* which refers to earnings changes for those with positive earnings at both $t-1$ and t (also calculable using

percentage changes), but also *labour market volatility*, i.e., including workers moving into or out of a job (not possible using percentage changes).

Having identified volatility at the individual level, the second (aggregation) step is to summarize the distribution of changes. The most commonly used measures are the standard deviation or variance of the distribution of arc percentage changes ($SD(\mathbf{a}_t)$, $Var(\mathbf{a}_t)$). Some researchers use the standard deviation of percentage changes, necessarily restricting attention to earnings volatility as defined above. Others use the standard deviation (or variance) of changes in log income, which approximates the standard deviation (variance) of percentage changes if all changes are small. Dispersion in the distribution of income changes might also be summarized using other measures such as quantile ratios.

Rather than summarizing dispersion in the distribution of income changes, a major strand of research has focused on the fraction experiencing a large income change between $t-1$ and t , especially the fraction with a large fall, on the grounds that these are likely to have the greatest adverse consequences for households. ‘Large’ is typically taken to be a change of at least 25%, sometimes at least 50%.

In this paper we mainly summarise two-period volatility in terms of the mean of the absolute values of the arc percentage change distribution, $Mn(|\mathbf{a}_t|) = (\frac{1}{N}) \sum_{i=1}^N |\mathbf{a}_{it}|$, where N is the total number of workers, a measure also used by Tomlinson (2018). This mean is equal to a weighted average of the income increases and decreases, where the weights are the fractions of individuals with increases and decreases. Jenkins (2025) shows that $Mn(|\mathbf{a}_t|)$ and $Var(\mathbf{a}_t)$ are related: the former summarizes changes using a linear (Manhattan or L1) distance metric and the latter uses a quadratic (Euclidean or L2) distance metric. This means that $SD(\mathbf{a}_t)$ and $Var(\mathbf{a}_t)$ give larger weights to relatively large income changes than $Mn(|\mathbf{a}_t|)$. Differences across indices are likely to be more apparent for estimates of labour market volatility than earnings volatility because the former includes arc percentage changes of $\pm 200\%$.

For all the volatility measures, repetition of the calculations over a set of moving windows of calendar times provides estimates of volatility trends.

Multiperiod measures come in non-parametric and parametric (model-based) flavours. Of the former, most measures have a common structure. First, calculate each person’s longitudinal mean income over $T_i > 1$ periods. Individual-level volatility is then represented by the income blips above and below this average and summarized by the dispersion of these deviations. (With annual data spanning many years, systematic variation in income with age

is typically removed via regression methods prior to calculations, or otherwise blips around the raw average would represent age-related mobility not only volatility. See, e.g., Latner (2018).) Overall volatility is the average of the individual-level dispersions.

Gottschalk and Moffitt (1994) pioneered this ‘window-averaging’ method. They measure individual earnings in logs and individual-level volatility is the variance of deviations from longitudinal mean log earnings. In contrast, Barnia and Leete (2009) and others since measure income in levels and summarize each person’s deviations from the longitudinal means using the coefficient of variation (CV) or its square.

Both of these variants of the window-averaging approach incorporate a Euclidean distance concept when summarizing income deviations from individual longitudinal means, just as $SD(\mathbf{a}_i)$ and $Var(\mathbf{a}_i)$ do in the period-pair context. Jenkins (2025) proposes a multiperiod counterpart to period-pair measure $Mn(|\mathbf{a}_i|)$ incorporating Manhattan distance. Define an individual-level income blip in period t , b_{it} , as the absolute difference from the longitudinal mean normalized by the longitudinal mean, $b_{it} = \left| \frac{y_{it} - \bar{y}_i}{\bar{y}_i} \right|$, and let individual-level volatility be the sum of the b_{it} over the T_i periods divided by the number of income changes ($T_i - 1$). The cross-person average of this normalised sum is the aggregate volatility measure and equals $Mn(|\mathbf{a}_i|)$ when $T = 2$. As for the period-pair case, the measure is less sensitive to large income changes than the averaged CV.

Another approach to multi-period volatility measurement is to classify individuals according to the shapes of their earnings trajectories, for example distinguishing between those that are stable, relatively stable with a few small blips, or erratic, with the last case being the most troublesome for families to manage. Brewer, Cominetti, and Jenkins (2025) compare volatility differences across groups using this method, inspired by the original trajectory classifications of Hills, McKnight, and Smithies (2006). See also this paper’s Appendix D. For yet another approach to multiperiod volatility measurement, see Bossert and D’Ambrosio (2011).

Multiperiod estimates of volatility and its trends can also be derived from econometric models decomposing inequality at a point in time into ‘permanent’ and ‘transitory’ variance components. The latter component is taken to represent volatility. In the most basic canonical model, the logarithm of income is equal to the sum of a time-invariant individual-specific component plus an uncorrelated component that is individual- and time-specific. Hence, inequality – measured by the variance of $\log(\text{income})$ – equals the sum of the variances of the permanent and transitory components. In the more sophisticated models now commonly

used, permanent components are allowed to change slowly over time, e.g., following a unit root process, and transitory components are assumed to follow some low-order autoregressive-moving average process. Income variances and covariances within and across multiple periods enable identification and estimation of the parameters describing the evolution of the variance components, and thence volatility.

Most applications are to employment earnings rather than household income. A pioneering study is Moffitt and Gottschalk (1995). Almost all researchers fit models measuring earnings in logs and inequality with the variance of logs, non-earners are excluded from analysis, and the focus is on men. One exception is Shiu, Zhang, and Gottschalk (2023), whose model uses income in levels, inequality is the coefficient of variation, and observations with no income from specific sources are included. Another exception is Ostrovsky (2012).

Whether researchers should use a parametric model-based approach or a non-parametric one to estimate volatility and its trends is a moot point. Variance components models provide an interpretation of how earnings inequality evolves by decomposing changes into (changes in) permanent and transitory variances. Neither approach directly identifies the fraction of income volatility that represents income risk of social concern, although combined with longitudinal data on consumption, a variance components approach can be used to estimate the extent to which households self-insure against transitory income shocks. See, e.g., Blundell, Preston, and Pistaferri who find “full insurance ... except among low-income households” (2008, p. 1887).

There are longstanding arguments that estimates of the variance components and their relative size are sensitive to model specification. The arguments are made most strongly by Shin and Solon (2009) who point out that Baker and Solon (2003), using a model that encompasses Moffitt and Gottschalk’s (1995) as a special case, reject the latter’s restrictions and derive different estimates of trends in the two variance components. See also Guvenen’s (2009) discussion of the difficulties of differentiating between different model specifications. These arguments about model specification could be usefully revisited: they may have lost force in recent years as relatively general models have become easier to fit (and data have improved).

If the focus is on volatility per se (rather than decomposing inequality changes), as it is here, the balance shifts in favour of the non-parametric approach for several reasons. First, although levels and trends in the non-parametric measures of volatility reflect permanent as well as transitory shocks, this is arguably a virtue rather than a problem:

The recent interest in volatility trends stems in large part from a concern about whether earnings risk has increased. Because permanent shocks ... are even more consequential than transitory ones, it makes good sense to include them in the measurement of earnings volatility. (Shin and Solon 2009, p. 9).

Second, Moffitt and Gottschalk (2012) have compared earnings volatility estimates from their variance components models and their non-parametric window-averaging approach and conclude that the two approaches provide similar conclusions about the pattern of trends over time in the transitory variance (for US men's earnings), although the latter does not identify exact turning points because it averages over years.

Third, relatedly, if general trends rather than specific turning points are of primary interest, the window-averaging approach "has the virtue of simplicity and transparency whose defects may not be qualitatively important" (Moffitt and Gottschalk, 2012, p. 218). This is a non-trivial advantage for communication of volatility research findings to policymakers and other non-specialist audiences. Winship makes a similar point even more strongly, stating that variance components models' estimates of volatility "measure a quantity that is a statistical construct and not observed or necessarily experienced by actual households from year to year" (2011, p. 12).

Fourth, non-parametric approaches are well suited to examinations of differences in volatility levels and trends across population subgroups. These are of particular interest to non-specialist audiences cited above. With variance components models, derivation of volatility breakdowns requires fitting the models to the different groups separately and that is rarely done. The non-parametric measures cited above are easily decomposed by population subgroup because they can be expressed as a weighted sum of subgroup measures, where the weights are the subgroup population shares. The analytical constraint is then simply the characteristics available in the dataset being used. Gottschalk and Moffitt (2009), armed with both model-based and window-averaging estimates, report breakdowns by education and position in longitudinal average income using the latter. Subgroup breakdowns are also presented in most of the research using non-parametric methods that we cite below.

3. Volatility in family and household income

3.1. Income volatility in the USA

The USA is the starting point for our overview of income volatility because it has received the most attention and is distinctive in other ways. Most studies have been based on a household panel survey, the Panel Study of Income Dynamics (PSID), but there are other longitudinal data sources that have also been used to triangulate findings and for robustness checking. By comparison with other countries, US income volatility – whichever measure is used – is at a high level and rose markedly between the start of the 1970s and the 2000s. There are few studies with findings for years after 2010.

To illustrate these points, we take as an initial reference point the summary evidence provided by Gottschalk and Moffitt's (2009) review article. The authors apply their window-averaging method to PSID data using nine-year rolling windows, providing estimates for 1974–2000. Income is the log of real annual pre-tax family income (excluding EITC and food stamps), adjusted for family size and composition using the relativities in Official Poverty lines. The sample is family heads aged 30–59. Gottschalk and Moffitt's (2009) Figure 5 shows a dramatic rise in over the period as a whole, increasing sharply during the latter half of the 1970s, flat in the early 1980s, but thereafter increasing steadily through to 2000. The overall increase in volatility is from around 0.08 to 0.14, i.e., around 75%.

Gosselin and Zimmerman (2008), using similar income definitions to Gottschalk and Moffitt but seven-year windows (alternate years only to accommodate the PSID's change to biennial interviewing in 1997) and some different sample selection choices including focusing on respondents aged 25–64, find broadly similar volatility patterns for most of the period. Their Figure 1 shows a doubling of volatility from around 0.12 at the start of the period rising to 0.25 in 1998. By contrast with Gottschalk and Moffitt, volatility is relatively flat in the mid- to late-1990s. Gosselin and Zimmerman's (2008, Figure 5) breakdowns by age provide an explanation for this contrast: volatility trends were similar by age group for most of the period, except in the mid- to late-1990s, when it was unchanged or fell for those aged 25–34 or 55+ (groups not fully covered by Gottschalk and Moffitt 2009). Gosselin and Zimmerman also report volatility being markedly higher for respondents from families in the poorest income fifth, and with education less than high school.

Nichols and Rehm (2014) show using the PSID that the rise in US income volatility also appears if window-averaging volatility is done using family income in levels (the averaged CV based on three-year windows). See their Figure 2 (measure V), which also

shows that the volatility rise continued over the 2000s. Their sample is persons aged 25–60 and they report that patterns are similar for both real pre-tax and post-tax (net) annual income.

Dynan, Elemendorf, and Sichel (2012) summarize volatility levels and trends using $SD(a_t)$ calculated using two-year changes to accommodate PSID interviewing frequency (and then averaged over 3-year rolling windows), and also the prevalence of large income changes. Their samples exclude retired people and students, and income refers to real annual pre-tax household income, not size-adjusted. $SD(a_t)$ volatility increased markedly between the 1970s and the late 2000s from around 0.38 to 0.50 (some 30%). The rate of increase was slower in the 2000s than before. The rise was driven by the growing prevalence of large income changes: the upward trend largely disappears if the top and bottom 25% of the arc percentage change distribution are excluded, and there is a rise throughout the period in the prevalence of income falls of at least 25% and falls of at least 50% (above trend in recessions, below trend in recoveries).

The ‘prevalence of a large income fall’ measure has been used extensively by Jacob Hacker and collaborators. For example, Hacker and Jacobs (2008) use PSID data to estimate the prevalence of falls between one year and the next in real pretax size adjusted annual family income among all individuals (not only adults) was around 3%–4% at the beginning of the 1970s but around 8% in the early 2000s, and counter-cyclical (as Dynan et al. 2012 report). The differences in prevalence estimates across studies (lower in Dynan et al.) are perhaps because of different income definitions and sample choices.

Hacker et al. (2014) use the prevalence of large income falls measure as the core input to their economic security index. The authors focus attention on matched data from successive years of the Current Population Survey (CPS) rather than the PSID. Their Figure 3 (‘income only’) series shows a rise in the fraction with a year-to-year fall in annual pre-tax size-adjusted household income of at 25% to rise from around 14% in 1986 to around 19% in 2010, an increase by one third, with counter-cyclical variation around the upward trend. Calculations are based on all individuals aged 18+.

The most comprehensive study of the period through 2010 using two-year matched CPS data is by Hardy and Ziliak (2014). Also notable is their use of a comprehensive disposable family income variable (with tax payments estimated using the TAXSIM model), not adjusted for differences in family size or composition. For a sample of household heads aged 25–60, Hardy and Ziliak report that $Var(a_t)$ doubled between 1980 and 2009 from around 0.12 to 0.25, with all the increase occurring before 2000/2001, after which volatility

was relatively flat. The metric matters: in terms of $SD(a_t)$, the rise was from around 0.34 to 0.50, an increase of 44%. This is larger than the increase in $SD(a_t)$ shown by Dynan et al. (2012, Figure 3), and may reflect use of a different income measure, dataset (PSID versus CPS) or sample selections, and two-year differences.

The four Survey of Income and Program Participation (SIPP) panels undertaken at intervals between the mid-1980s and early 2000s have also been used to study volatility. Gosselin and Zimmerman (2008, Figure 9) report that the rise in their window-averaging volatility measure for family income is substantially smaller according to the SIPP than the PSID. Winship (2011) and Hacker et al. (2014) compare estimates based on the SIPP, two-year CPS, and PSID, with both studies also including extensive discussions of issues such as attrition and how to treat imputations for missing incomes. (See also Dynan et al., 2012.) Winship argues that imputed observations should be excluded, in which case he finds all three data sources provide similar estimates of the level and trend in the prevalence of large income falls of at least 25% (compare his Figures 4a and 4c). Compared to PSID-based estimates such as those of Hacker and collaborators, Winship's estimates show a smaller rise over time. In contrast, Hacker et al. (2014, Figure 6) report that their PSID data produce higher level estimates than the other sources and they also point to an upward trend in all estimates regardless of dataset.

Dahl, DeLeire, and Schwabish (2011) also use SIPP data, comparing estimates derived using survey-reported incomes and an income measure in which survey-reported employment earnings are replaced by Social Security Administration (SSA) administrative records on earnings for linked respondents. In either case, income is real annual household pre-tax income, and estimates are derived for household heads aged 25–55. Dahl et al. point out that use of linked administrative earnings data may address issues of missing data and imputation (a non-trivial issue in the SIPP as with other panels), but they also acknowledge that linkage brings its own problems. (They report 10% – 20% of SIPP respondents could not be matched in 7 panels, and around 40% in the 2001 panel.) Using the SIPP-SSA linked data, the fraction with an income fall of at least 50%, and fraction with a rise of at least 50%, are each about 4%–5% between the early-1990s and mid-2000s (Figure 1). $SD(a_t)$ is unchanged over the same period as well, at around 0.3 (Figure 2). In terms of volatility levels, the prevalence of large income falls is greater for heads in the poorest fifth of the household income distribution, those with less than high school education, or aged 25–29, thus echoing Gosselin and Zimmerman's (2008) findings cited above. Dahl et al. (2011) also compare volatility estimates based on survey-reported household income for samples including and

excluding imputations, showing that their exclusion implied estimates close to linked SSA data estimates.

Debacker et al.'s (2013) study is the only US paper about income volatility that is entirely based on administrative record data (IRS tax return data). This brings strengths such as very large samples, no imputations or top-coding, and less measurement error. The authors' panel of IRS tax return data spans 23 years from 1987 to 2009, and they construct relatively comprehensive pre- and post-tax income definitions (though self-employment income is not included) and analyse tax filers aged 25–60. Analysis is of residuals from regressions of household income on sex, age, and household composition. Remarkably, there is virtually no change between 1988 and 2009 in the SD of percentage changes in pre-tax household income (whether 1-year or 2-year changes: see their Figure 6). No trend is found using a Gottschalk-Moffitt window-averaging measure either (Debacker et al., 2013, p. 101). The authors also fit variance component models to their household income data and samples. Their transitory variance estimates (Figure 8) are consistent with the non-parametric ones in the sense of showing virtually no rise over the post-1990 period.

Clearly, survey and administrative data estimates provide quite different pictures of trends in household income volatility in the USA – rising since the 1980s according to the former but not according to the latter. In addition, there are no published estimates for the period after around 2010, more than a decade ago, even though the PSID, CPS, and SIPP have continued. What has happened to household income volatility more recently?

Having access to suitable up to date administrative datasets would also help address this issue and allow further checking of the different estimates of trends cited above. Relatedly, the US Census Bureau's MOVS project (Jones et al., 2024) innovatively combines demographic and IRS tax administration records to build a new longitudinal data set with household income data for the population of working-age adults. The downloadable [Data Tables](#) provide new national estimates of the fractions with an increase in income of at least 25% and of a fall of least 25%, and $\text{Var}(a_i)$, for one-year changes over the period 2006–2019. $\text{Var}(a_i)$ remains fairly constant, at around 0.37 between 2005 and 2017, before falling to 0.34 in 2018, but rising sharply to 0.42 in 2019. The fractions with a large fall and a large rise remain constant over the 2010s at around 16% and 20%. The levelling off in income volatility in the 2000s suggested by earlier studies appears to have continued, albeit with a large uptick in the last year before Covid-19 onset.

In sum, most US research finds a substantial rise in volatility the 1970s and 1980s with a slowdown thereafter especially from the 2000s onwards, especially according to

research based on administrative data. There are variations in the headline story depending on the choice of volatility measure, and the different ways in which authors treat imputed or top-coded observations, and define their sample selection criteria. Almost all studies take an annual perspective, and few provide information about the post-2010 period.

3.2. Income volatility in non-US countries

The income volatility experiences of countries other than the USA is strikingly different in terms of both levels and trends. (Non-US estimates are all based on household panel surveys, first available for West Germany starting 1984 and Britain from 1991.) For example, Jenkins (2011) shows using British Household Panel Survey (BHPS) data that household net income volatility in the UK was constant or declined slightly between the 1990s and mid-2000s, regardless of whether the measure is a Gottschalk-Moffitt window-averaged transitory variance, the fraction of persons with a large income fall, or the standard deviation of two-year log income changes. Avram et al. (2022) using the BHPS's successor (the UK Household Longitudinal Survey, UKHLS) show that volatility according to $SD(a_i)$ had a shallow U-shape between 2010 and 2017, starting at around 0.42, declining a few percentage points and later returning to the same level.

Comparing volatility levels, the window-averaged transitory variance for Britain in the early 2000s was around 0.04 (Jenkins, 2011, Figure 6.1b) but around 0.14 in the USA (Gottschalk and Moffitt 2009, Figure 5). At the same time, the probability of a one-year income fall of at least 50% was only around 3% in Britain (Jenkins, 2011, Figure 6.6) but around 7% to 8% in the USA (Hacker and Jacobs 2008, Figure C). Focusing on Britain and Germany and also using a window-averaging measure, Bartels and Bönke (2013) confirm the constant volatility finding for the period from mid-1980s to mid-2000s. Menta et al. (2021), also using a window-averaging measure, showed that volatility in Italy approximately doubled between the mid-2000s and mid-2010s, but the increase was much less than that in the USA and the 2014 level was less than half the US level.

The exceptional experience of the USA is further highlighted by comparisons involving more countries. For example, Nichols and Rehm's (2014, Figure 2) window-averaged measure V remained at around 0.01 from the mid-1990s to the late 2000s for Britain, Canada, and Germany, whereas the US estimate rose from around 0.02 to more than 0.03. Hacker and Rehm (2009) report larger income-fall prevalence estimates for many countries. Focusing on the most comparable estimates, for Britain, Germany, and the USA, they find the prevalence of income drops to be highest for the USA at every threshold. More

recently, Hacker (2018) reports estimates for 35 countries of the share of persons with a one-year fall in size-adjusted real household income of at least 25%, showing the mean and range of share estimates over the period for which each country has estimates (the latest estimates refer to 2014). Mean shares range from around 6% for Sweden and Norway to Spain at around 20%, Serbia at 25%, with the estimates for the USA around 17%, for Australia and Britain 13%, and Germany and France 10%. The US's relatively high fraction with a large income fall is confirmed in Hacker and Rehm's updated estimates for 21 countries (2022, Figure 1).

3.3. Decomposing income volatility by income source

A substantial advantage of household panel survey data over administrative record data is that the former has much more detailed information about the various income components comprising total household income and household composition. This facilitates analysis of the extent to which households and welfare states (income taxes and cash transfers) offset the volatility arising from market incomes, especially labour earnings, and how this may differ across countries or be changing over time within a country.

There is of course much comparative literature about welfare states as providers of social insurance but Hacker and Rehm (2022) argue that specific measures of welfare states' insurance impact are relatively rare. They develop the concept of risk reduction, i.e., "the degree to which taxes and transfers reduce the frequency and severity of major losses in household income" (2022, p. 456), summarized by the difference between the share of adults experiencing a large fall in market income and the share with a large fall in disposable income (attributing the differences to taxes and transfers). Hacker and Rehm find, perhaps unsurprisingly, that risk reduction is greatest for Nordic countries (around two-thirds for Norway) and relatively low for Anglo countries like Britain and the USA at around 30% (2022, Figure 4).

Volatility trends within countries have also been studied exploiting detailed information about household income packages. A leading early example is Dynan et al. (2012) for the USA, showing inter alia how declining spousal earnings volatility offset the impact of rising head's earnings volatility on household income volatility, but rising capital income and transfers volatility had a reinforcing effect. (Volatility is measured using $SD(a_t)$ and so zero values for component incomes can be included.) Gottschalk and Moffitt (2009) come to similar conclusions. Although using a window-averaging measure based on log

incomes, they avoid dealing with zero incomes by aggregating components so that the prevalence of zeros for each sub-aggregate is negligible.

The most comprehensive decomposition of trends in US family income volatility is by Hardy and Ziliak (2014). They exploit the property that $\text{Var}(a_t)$ for total income for a given year can be expressed in terms of a sum of the weighted volatilities of each income source plus the covariances between weighted source volatilities, where the weight on each source volatility is its share in total income. For their analysis of levels and trends between 1980 and 2009, Hardy and Ziliak distinguish five components of family disposable income for heads aged 25–60: head's earnings, spouse's earnings (if spouse is present), non-transfer other income, transfer income, less net tax payments (which include, e.g., tax credits such as EITC). Hence, changes over time are related to changes in variances, covariances, and income shares (around 20 terms). The authors' summary of the complex pattern of changes over the three-decade period is that the “increased volatility comes from higher instability of head and spouse earnings, and other nonlabor income, as well as from a reduced covariance between these income sources with the tax system. This suggests that current tax policy is less effective in mitigating income shocks than previous decades” (Hardy and Ziliak, 2014, p. 459). Using shift-share analysis, Hardy and Ziliak (2014) show that changes in component income shares also played a role, but family income volatility would have risen even if the shares had remained fixed at their 1981 values.

Avram et al. (2022) used Hardy and Zilak's (2014) decomposition method to assess the trends in UK household income volatility over the period 2009–2017, considering working-age individuals (aged 25–59) and older individuals (aged 60+) separately. As found in other studies, household earnings volatility is lower than head's earnings inequality (including spousal earnings lowers volatility), and taxes and cash transfers reduce volatility still further. Echoing Hardy and Ziliak (2014), Avram et al. (2022) also conclude that the ability of taxes and transfers to offset labour income shocks declined, referring to decreases over time in covariances between changes in individual earnings on the one hand and changes in benefit incomes and changes in tax payments on the other. But, at the same time, the volatilities of earnings and non-labour income fell over the period for working-age individuals providing an offsetting force, thereby accounting for the shallow-U shaped trend in household net income volatility over the period.

3.4. Sub-annual perspectives on income volatility derived from surveys and small-N studies

All the studies reviewed so far are distinguished by their annual perspective, considering changes in annual income or changes in monthly incomes a year apart. Few studies have examined household income volatility within a year. Each four-monthly interview wave of the US SIPP collects monthly income data over the previous four months, but these are rarely used because of between-wave ‘seam bias’ problems (correlations of within-wave incomes noticeably greater than for between-wave incomes). Hence, for example, Acs, Loprest, and Nichols (2009) focus on four-month incomes. (Cf. Western et al. 2016 who use the first monthly income in each SIPP wave.) Acs et al. (2009) examine the prevalence of large income drops between waves, drawing attention to their relatively high prevalence and variation by income group (highest for the top and bottom fifths on average). The authors also document how, although such drops are temporary for a significant minority (almost 40% return to their pre-drop income within a year), the drop is persistent for many too (for more than 20%, income is less than half the pre-drop level for at least a year). Acs et al. comment that “[e]ven a short-term substantial drop in income may adversely affect these families due to the attendant stress and loss of resources associated with such sudden declines”. (2009, p. 8).

Bania and Leete (2009) is one study that does use the SIPP’s monthly data to examine pre-tax family income volatility. Their measure is the median of the coefficient of variation (CV) of deviations from an individual’s 12-month longitudinal average pre-income. Comparing estimates from the 1991/2 and 2001 panels, the authors document that volatility was substantially higher for low-income households, and that it increased substantially for this group between 1992 and 2003 (Bania and Leete, 2009, Table 1), which they attribute to a shift in the composition of income from welfare payments (relatively stable) to earnings and other income sources (more unstable). The authors discuss the SIPP seam bias problem in detail and reanalyse their data using two separate adjustments to address it (including using only one month’s income per wave). Their conclusion is that their main findings are little affected by the seam bias issue. A later study using additional SIPP panels to study within-year volatility, but focusing on children, is Morris et al. (2015).

A recent study of within-year income volatility is OECD (2023). This presents estimates of monthly volatility over a 48-month window between 2013 and 2018 for 20 European OECD countries derived from EU Statistics on Income and Living Conditions panels, using samples of households with stable composition over the four years and whose main employment income earner is aged 18–59. Income is a gross income concept (market

incomes plus unemployment benefits, old-age pensions, and educational allowances), adjusted for differences in household size. Volatility is summarized using the longitudinal (48-month) squared CV averaged across households. OECD (2023, Figure 1.2) document substantial cross-national differences in monthly volatility. Four countries have estimates of around 0.05 or slightly less (Czechia, Norway, Slovakia, and Luxembourg), six countries' estimates lie between 0.05 and 0.10, eight have estimates between 0.10 and 0.15, Spain's estimate is around 0.20, whereas the UK's estimate is around 0.28.

An important caveat is that EU SILC panels collect annual income information at each of the four annual interviews, not monthly income data. OECD (2023) constructed monthly income measures exploiting the monthly employment histories that are available, allocating reported annual incomes pro rata to the months in which individuals were in full- or part-time employment. Private pension income was allocated to months in which individuals were retired or unemployed, and income from other sources was spread evenly across the year. Although the cross-national patterns in Figure 1.2 are plausible, the UK's outlier position raises questions about the construction method or at least about differences in how SILC is implemented in the different countries. For example, the UK survey differs from other countries in collecting information about 'current' income (round about the time of the annual interview), not annual income as for other SILC countries.

Sub-annual household income volatility has also been examined using small-*N* studies of family finances combining granular quantitative and qualitative data collection. The leading UK example is by Hills et al. (2006). The authors had complete information on weekly incomes from all sources over one calendar year for a sample of 93 low- to middle-income working families with children, all receiving Working Families Tax Credit (WFTC), a precursor of today's Universal Credit, in the Winter of 2002/03. WFTC made up 24% of couple households' income and 30% of lone parents' income. The incomes for financial year (FY) 2002/03 were grouped into 13 four-weekly periods for analysis.

Hills et al. (2006) report that although a few families had stable disposable incomes over the year, many experienced substantial instability. For example, using the averaged CV measure, volatility for all 93 cases was 16.5%, and a third had CVs greater than 20%. (In contrast, Bania and Leete (2009) report, for US families in the 2001 SIPP panel, median CV estimates of 32% with those with household incomes between 50% and 100% of the official poverty line and 28% for those with incomes between 100% and 150% of the official poverty line.) Moreover, "[a] quarter of the cases had 'erratic' or 'highly erratic' reported incomes, with at least four of the thirteen periods outside the range from 85 to 115 per cent of their

annual average” (Hills et al. 2006, p. 4). Greater instability was experienced by families with low incomes, lone parents, and a few who started the year without labour earnings. Hills et al. (2006) report that all the major income components – labour earnings, other income (including child support), cash transfers and WFTC receipts – were all highly variable but, whereas adding in other market income increased variability compared to net pay, adding in either social security benefits or tax credits reduced variability, an effect that was greatest for the most disadvantaged groups (e.g., lone parents, tenants). In contrast to our findings reported in Section 5, Hills et al. also state that “there was no pronounced seasonality in income receipts over the year. There was somewhat higher net pay in the period just before Christmas” (2006, p. 30).

The leading US ‘family finances’ study is the Financial Diaries project, collecting comprehensive data on cash inflows and outflows for 235 low- and moderate-income households for around 12 months between 2012 and 2013. For details, see e.g. Hannagan and Morduch (2016), Morduch and Siwicky (2017), and Morduch and Schneider (2018). The authors paint a picture of marked fluctuations in incomes over the year, similar to the UK situation portrayed by Hills et al. (2006). For example, “[h]ouseholds experienced, on average, 2.7 spikes (months with income above their monthly average by at least 25 percent) and 2.7 dips (months with income below average by at least 25 percent)” (Hannagan and Morduch 2016, p. 238), and greater for low-income households, with correspondingly greater experience of within-year poverty spells. The median of the CV volatility measure was around 34% and the mean around 39% (Hannagan and Morduch 2016, p. 247), i.e., of roughly the same magnitude as reported by Bania and Leete (2009).

The Financial Diaries study reports that spending tended to track income to a greater extent for low-income families than for more well-off families – they are less able to smooth consumption. Hills et al. (2006) did not track spending, but their qualitative interviews provide evidence consistent with the US finding: respondents referred to adapting their spending to their incomes, trying to live within their means, and facing problems if there were unexpected outgoings. More recent UK family finance studies such as Angsten Clark and Otulana (2024), Biosca et al. (2020), and Griffiths and Wood (2024), provide additional evidence of these problems.

4. Volatility in individual employment earnings

Research on earnings volatility is much more extensive than on family and household income volatility. Earnings volatility is of particular interest to labour economists, as a topic in its own right (as a signal of labour market ‘risk’), and also related to their interest in the dynamics of earnings inequality – hence also more widespread use of technically sophisticated models to estimate (trends in) permanent and transitory earnings variances. There are new and more up-to-date data sources about earnings compared to income, notably administrative record data from private and government sources. Although most studies take an annual perspective, a few have exploited new administrative sources to examine within-year volatility. Another development is that earnings volatility estimates are increasingly produced for women as well as men.

Our review of what we know about earnings volatility focuses on the more recent literature. The review of income volatility in the previous section has already indicated the main findings for the period up to the mid-2010s: the principal driver of income volatility levels and trends was the volatility of family head’s labour earnings. Moreover, extensive reviews of that period already exist, including Gottschalk and Moffitt (2009), Moffitt and Gottschalk (2012), Jenkins (2011, chapter 6), and the introductory sections to later articles that we cite below. Analysis is of individuals’ gross earnings, i.e., earnings prior to the deduction of taxes and social insurance contributions.

4.1. Earnings and labour market volatility in the USA

A useful stepping-off point is the US situation portrayed by Moffitt and Gottschalk (2012), focusing on male household heads aged 30–59, using PSID data covering 1970–2004. Using three methods (variance components models, window-averaging, and another non-parametric approach), they show that volatility doubled between the start of the 1970s and the mid-1980s and was roughly constant until the end of their data period. There is also cyclical variation from trends, with volatility higher after recessions.

Shin and Solon (2011) also use PSID data, but for male heads aged 25–59 over the period 1970–2006, and they measure volatility using the SD of log earnings changes. Their headline findings about volatility levels and trends are broadly the same as Moffitt and Gottschalk’s (2012), except that Shin and Solon estimate a more distinct fall in volatility during the 1990s with a rise again during the 2000s following the recession at the end of the 1990s (2011, Figure 1). In contrast, Sabelhaus and Song (2011) using administrative data, the

Social Security Administration Continuous Work History Sample (CWHHS) one percent Master Earnings File (MEF) sample, report a more continuous decline in volatility (variance of log earnings changes) from the mid-1980s through to the mid-2000s. Although Sabelhaus and Song's samples consist of both men and women, they report that similar temporal patterns hold for men and women separately (2011, fn. 13), which suggests the different estimates of trends across studies are related to the use of different data sources.

Relatedly, consider the analysis of Ziliak et al. (2011) based on matched CPS panels, also distinctive for providing separate estimates for men and women and measuring individual-level earnings changes using the arc percentage so that non-positive earnings observations could be included (unlike in the three US studies cited in the previous paragraph). Ziliak et al. (2011, Figures 2 and 3) show that if their samples are restricted to wage and salary earnings with positive earnings, trends in volatility are similar – relatively flat for men over the two decades up to 2008 but declining for women – regardless of whether they use their $SD(a_t)$ measure, Shin and Solon's, or Gottschalk-Moffitt's window-average measures. When Ziliak et al. (2011) include observations with zero earnings (labour market volatility), the trends over those two decades are much the same for women, but volatility levels are markedly higher, as expected. Decomposition analysis leads the authors to conclude that “men's earnings volatility is increasingly accounted for by employment transitions, especially exits, while the share of women's volatility accounted for by continuous workers rose” (2011, p. 742). However, the relatively flat level of volatility in the 2000s for men estimated from the matched-CPS and PSID data remains at odds with the more distinct fall reported by Sabelhaus and Song for the same period, using administrative data.

The question of whether the differences in estimates of earnings volatility trends and levels are due to using different data sources was recently addressed head-on. Moffitt et al. (2023) review data-related sources of potential discrepancies in past research, discuss the strengths and weaknesses of the sources, and provide new comparable and more up to date estimates using them. To isolate the impact of different data sources, the analysis used the same volatility measure, $Var(a_t)$, and applied the same selection criteria, restricting attention to men aged 25–59 with positive earnings for years $t-1$ and t . Three series of estimates are entirely panel survey-based (PSID, SIPP, and matched-CPS, each using survey responses on earnings), two are survey-based but with administrative data on earnings (matched-CPS and SIPP using linked Social Security Administration data), and one source is entirely administrative (the Longitudinal Employer-Household Dynamics (LEHD) dataset which

includes earnings for almost all workers covered by unemployment insurance in states that provided data). Only the PSID data covers the 1970s and it is only from the late-1990s onwards that there are estimates for all six sources.

Moffitt et al.'s (2023) Figure 1 shows the baseline comparison of volatility levels and trends. The central finding is that “when put on a comparable basis, male earnings volatility in six survey and administrative datasets shows no sign of a major net increase or decrease since the late 1980s or early 1990s although experiencing significant countercyclicality” (Moffitt et al., 2023, p. 8). This finding is robust to using other volatility measures (e.g., variance of log earnings changes), inclusion of zero values, and to using age-adjusted earnings rather than observed real earnings. The team also show that different types of trimming of the highest and lowest earnings in each cross-section can lead to substantive differences and, relatedly, point out that some of the differences between the LEHD series and the others is explained by the former's more extensive coverage of the bottom of the earnings distribution.

Blundell et al. (2024) take this research forward in several ways. They use the same dataset as employed by Ziliak et al. (2022) in the Moffitt et al. (2023) project – CPS respondents linked to their Social Security Administration earnings records – but with the linked earnings records extending further back in time (to 1978) and further forward (to 2019). The authors also derive separate volatility series for six groups (sex by race by education), as well as exploring differences by age and birth cohort. Blundell et al.'s Figure 2 summarizes level and trends in $\text{Var}(a_t)$, showing for all groups a marked decline in volatility during the 1980s (especially black men with at least some college education), but broadly flat from the mid-1990s onwards apart from a distinct rise following the Great Recession followed by a fall back. The estimates reveal clear volatility penalties for being a black man and not having a college education. By 2019, volatility for black men with ‘some college education or less’ was more than twice that for white men with ‘college education or more’. For women, volatility also declined over the period as a whole, but more sharply initially from a higher 1980s baseline before flattening out in the 1990s. The Great Recession volatility uptick is less marked for women than men, and volatility declines at the end of the 2010s. Education and race group levels and trends also differ from men. For example, for most of the period, it is college-educated black women who have the lowest volatility. By 2019, there are no differences between black and white women: what matters is whether they have completed a college education. Blundell et al.'s (2023) Figure 3 shows that volatility trends are broadly the same for men and for women if calculations are restricted to those with

positive earnings at $t-1$ and t . Blundell et al. (2023) also fit various variance component models and decompose earnings inequality levels and trends into contributions from permanent and transitory variance. For all six groups, there is a downward trend in year-on-year transitory variance which the authors associate with a reduction in volatility.

4.2. Earnings and labour market volatility in non-US countries

Cross-national comparisons of earnings volatility with the US experience have been difficult because different studies use different datasets, and different coverage in term of time periods and sample selection (e.g., by age and sex). But, in so far as one can tell and as for income volatility, there is substantial heterogeneity across countries. For example, using methods as comparable as possible with Ziliak et al. (2011), Cappellari and Jenkins (2014, Figure 3) show that earnings volatility for British men and for British women was constant between 1992 and 2008, as it was for the USA according to Ziliak et al.'s estimates, but volatility levels were substantially higher in the USA. By 2008, $\text{Var}(a_t)$ was around two-thirds higher for US men, and around fifty percent higher for US women. Another transatlantic contrast was that labour market volatility for British men and women declined over the period but showed no trend according to Ziliak et al.'s (2011) US estimates. We update Cappellari and Jenkins's (2014) panel survey-based evidence for the UK in Section 5 and show that labour market volatility continued to decline while earnings volatility remained constant.

Cross-national heterogeneity is also revealed by Beach et al.'s (2014) comparisons of estimates of Canadian earnings volatility over the period from mid-1980s through to the early 2000s with estimates for the USA using Gottschalk and Moffitt's (2009) window-averaging volatility measure. Beach et al. remark that "while the transitory variance for women in the U.S. showed no major evidence of an overall trend over the 1990s, [their] study for Canada finds a broad decline in transitory variance over the 1990s followed by a marked increase in the 2000s" (2014, p. 585). Note also, in the light of the earlier discussion about types of data, that Beach et al. (2014) use administrative data, the Longitudinal Administrative Database (LAD), comprising earnings reported on tax forms of a 20% random sample of tax filers, whereas Gottschalk and Moffitt (2009) used the PSID survey-based responses.

The recent multi-country GRID project addresses comparability issues head on. All 13 country teams use administrative record data of some form, all consider real gross employment earnings from all jobs, and each country paper in the 2022 *Quantitative Economics* GRID special issue provides estimates of earnings and earnings change distributions and their trends derived according to a common analytical template and focuses

on samples of men and women aged 25–55. All 13 countries have estimates for the period covering the mid-1990s through to the late-2010s (but pre-Covid), and many have estimates going further back in time. In addition to the USA and UK, the countries include Canada, seven European countries (Denmark, France, Germany, Italy, Norway, Spain, Sweden) and three Latin American ones (Argentina, Brazil, Mexico), so the coverage is of a wide range of different types of economy and labour market.

Güvenen et al.'s (2022) editorial introduction to the GRID studies provides explicit cross-national comparisons of earnings volatility in terms of one measure, the variance of one-year log earnings changes (except for the UK where estimates refer to the $p90 - p10$ quantile difference), and separately for men and women. Thus, the focus is on labour market volatility. They take no account of differences in transitions into and out of jobs across these economies, and there is likely to be substantial variation in this, e.g., higher in Latin American countries. Moreover, the administrative data for the three LAC countries refer to workers in the formal sector, another factor hindering comparability, even though each of the three country studies also deploys household survey data to elaborate their analysis.

Güvenen et al.'s (2022) Figure 7 summarizes volatility levels and trends for the 13 countries. The authors state that this “paints a somewhat mixed picture with volatility flat for about half of the countries, declining for some countries and rising for others. It does not provide any evidence of a widespread rise in volatility or income risk around the world” (2022, p. 1342). Cross-national heterogeneity in volatility levels is also apparent. The USA is one of the countries for which volatility is shown as broadly flat for both men and women, with a value of 0.55 for both sexes. Few countries have higher volatility levels. There are the three Latin American countries – though for Argentina and Brazil, levels for men and women converge to the US level at the end of the period – and perhaps surprisingly Norwegian men (similar levels as Mexican men and women, 0.65). For most countries including the USA, men and women experience much the same volatility levels, with Norway and Sweden being outliers, with men's volatility distinctly higher than women's and it is unclear why. Germany and France exhibit the lowest volatility levels, for men and women (between 0.35 and 0.45). All in all, among the high-income countries, US earnings volatility levels stand out as high.

As we have shown, the strengths of administrative record data in terms of reliability and large sample sizes are increasingly being utilized. But it remains the case that almost all earnings and income volatility research takes an annual income perspective which hides volatility within a year. In contrast, our research reported in the next section focuses on within-year earnings volatility. Before discussing our findings, based on a new UK

administrative data source, we refer to the few other studies we are aware of that have also examined monthly earnings and income volatility using administrative record data.

4.3. Using administrative record data to study within-year volatility

Monthly transaction data for bank account holders have been used in three studies of month-to-month volatility, two for the USA, one for the UK. Farrell and Grieg (2015) leverage JPM Chase bank data for a sample of 100,000 primary account holders from their customer base to assess volatility in incomes and spending between October 2012 and December 2014. Farrell et al. (2019) do more: the sample size is increased to six million JPMC primary account holders, the time frame is extended, from October 2012 to December 2018, and there is more extensive analysis, including use of volatility measures more commonly used by academic researchers (such as the averaged longitudinal CV). Farrell and Grieg observe that their “sample is biased in favor of individuals who earn and spend more than the average individual nationally” (2015, p. 24), largely because they restrict attention to account holders with at least \$500 (2015 study) or at least \$400 (2019 study).

The headline finding from the 2015 report is that there is a large amount of month-to-month volatility. The 2019 report underscores this, stating that although income instability remained relatively constant between 2013 and 2018, “[t]he level of income volatility remained high, with those at the median level experiencing a 36 percent change in income month-to-month. ... There is wide variation in the levels of income volatility families experience, and volatility is greatest amongst the young and those in the highest income quintile” (Farrell et al., 2019, p. 2). To benchmark their data against other estimates, Farrell et al. (2019, Figure 3) also calculate the SD of annual log labour income changes for male account holders aged 20–64. Their estimate hardly changes across the six years, consistent with the lack of trend shown by the US LEHD data in GRID (Güvenen et al., 2022) but rather smaller in terms of levels (around 0.3 compared to the GRID estimate of 0.55). Part of the difference is because the Farrell et al. (2015, 2019) data refer to take-home (post-tax) income, not gross earnings as in GRID and other studies.

Farrell et al. (2019) find that some seasonal patterns in their data, mostly attributed to the incidence of December pay bonuses and months including five Fridays (hence more pay payments than other months) but conclude that calendar effects are relatively unimportant. (Comparisons of estimates adjusting for these effects and those that do not are much the same.) We examine seasonal patterns in detail in our UK research reported below.

For the UK, Tomlinson (2018) uses transaction data from over seven million Lloyds Bank Group (LBG) accounts covering financial year 2016/17, and focuses on the volatility of ‘pay’, i.e., total take-home employment earnings of individuals (aged 18+). Lloyds is among the top five of UK banks, with national coverage. One limitation of the data is that pay data could not be determined for joint accounts, potentially undermining the representativeness of partnered individuals.

A headline finding is that “pay fluctuations are the norm for the majority of employees. Only 9 per cent of employees who remained with the same employer throughout 2016-17 had no months in which take-home pay changed by a notable amount (greater than five per cent, either up or down)” (Tomlinson, 2018, p. 5). Moreover, there is a U-shaped relationship between pay level and volatility for those with same employer all year: “the absolute average monthly pay change for those with a steady job is highest for those on the very lowest earnings ... and lowest for those with annual take-home pay close to the median” (2018, p. 7). Tomlinson (2018) also reports that men and women experienced similar monthly pay volatility on average, but low paid men experienced distinctly greater volatility than low paid women. In addition, monthly pay was most volatile for young workers compared to middle-age and older workers.

The only sub-annual income volatility study based on UK government administrative record data (besides the current paper) that we are aware of is by Johnson et al. (2025). They use monthly data about 70,000 households receiving Universal Credit across seven local authorities from April 2022 to March 2023. (Universal Credit is the main means-tested cash transfer in the UK, available to working and non-working households.) The study documents widespread volatility income: for example, 30% of all households (36% of working households and 21% of non-earning households) experienced ‘erratic’ or ‘highly erratic’ changes in monthly income, and 21% of all households experienced ‘highly erratic’ changes. (The classifications are based on those of Hills et al., 2006.) This is innovative research, albeit limited by its non-representativeness of the population by income and geography.

Administrative records from a US payroll processing firm are the source of the data used by Ganong et al. (2024) to study monthly earnings volatility. Most of the payroll processor company’s clients are small firms though Ganong et al. argue that their data “appears to be representative of the prices of labor, the quantities of labor, and contract types for U.S. workers overall” (2024, p. 6). Their main samples cover between 10 and 25 million workers, covering 2010 through 2023. Although Ganong et al. mostly work with 1% or 5% samples for computational tractability, the samples are clearly very large. A particular

advantage of their dataset is that they can distinguish between base pay and additional pay (if present) such as overtime, paid holidays and vacations, tips, commissions, and bonuses. They also know work hours, whether hourly-paid or salaried, age, sex, and number of dependent children, and employer's industry.

Ganong et al.'s "main finding" is that "workers face substantial earnings fluctuations from month to month ... earnings changes are ubiquitous. In almost three-quarters of months, the change in pay is at least 17%" (2024, p. 8). Importantly, Ganong et al. also document that the principal source of the volatility is pay items other than base pay: "the standard deviation of monthly earnings changes is 28%, while the standard deviation of base wage changes is only 2%" (2024, p. i). Volatility is also higher for hourly paid workers compared to salaried ones. Their take-away is that "high-frequency labor market shocks are an important source of risk and fragility which has been masked by past studies of annual earnings" (2024, p. i). Our research reported below provides further 'unmasking' of earnings volatility in the UK context using a nationally representative dataset, though unfortunately our data do not have as rich a set of worker and firm characteristics as Ganong et al.'s.

Nationally representative data from Danish administrative registers covering January 2011 to December 2018 are used by Druedahl et al. (2023) to analyse changes in monthly earnings from all jobs, focusing on a sample of 400,000 never self-employed men aged 35–60. The original data sources are the reports that Danish employers are mandated to provide to the national tax agency (similar to our PAYE dataset discussed below). The principal goal of Druedahl et al.'s (2023) paper is to develop a permanent-transitory variance components model of earnings allowing for infrequent shocks that can account for the principal features of both monthly and more commonly used annual data. However, they do report some descriptive features of their monthly earnings change distributions, emphasizing the high proportion of employees with no log earnings change from month to month (approximately one half). This is a relatively high fraction and perhaps a Danish phenomenon – recall from the earlier GRID project discussion that Denmark has comparatively low (annual) earnings volatility.

The stage is now set for us to discuss our research based on a new UK administrative record dataset on earnings, providing a new look at pay volatility from a high frequency perspective.

5. Earnings volatility in the UK: new evidence from a high frequency perspective

In this section, we provide new evidence about the volatility of employment earnings (pay) in the UK, monthly over five financial years. Ours is the first UK study to use governmental administrative data to analyse this topic, and we benefit from the nationally representative nature of the data, the very large samples, and lower measurement error than in survey data.

Tomlinson (2018) provided the first UK evidence about monthly pay volatility but for only one year (2016/17), and our data are more representative of the UK employee population than the bank account transaction data he used. We also document seasonality more extensively and provide multiperiod as well as period-pair volatility measures.

As background, we begin by providing updated earnings and labour market volatility estimates from an annual perspective (monthly earnings a year apart) using three UK panel survey datasets. Next, we turn to our administrative data. After explaining their nature, we set out the main features of month-to-month earnings volatility in the UK, drawing attention to its distinct seasonal patterns, and how volatility and seasonal patterns differ across subgroups, especially those at the top and bottom of the earnings distribution. We also examine volatility from a multiperiod perspective, summarizing the instability of pay over the 12 months of a financial year.

5.1. Updated survey-based volatility estimates from an annual perspective

There are currently three UK panel surveys that can be used to study earnings and labour market volatility, each providing an annual perspective. There are household surveys that can provide estimates of both labour market and earnings volatility – the British Household Panel Survey (BHPS) and its successor the UK Household Longitudinal Study (UKHLS), and the quarterly Labour Force Survey (LFS). In addition, there is the Annual Survey of Hours and Earnings (ASHE), a large annual survey of employees with reporting by employers that can be linked longitudinally and provide earnings volatility estimates.

Almost all previous research about UK earnings volatility reports that it has hardly changed since the early 1990s. Tomlinson (2018, Figure 1) compares estimates based on the LFS for 1998–2017, with estimates from the BHPS (used by Cappellari and Jenkins 2014, covering 1991–2008) and its successor the UKHLS (which began in 2009). Pooling men and women, Tomlinson reports that the BHPS/UKHLS and LFS provide similar estimates of earnings volatility: $SD(a_t)$ is around 0.30 across the period. He does not report labour market volatility estimates; Cappellari and Jenkins (2014, BHPS data) report that this was declining

for men and women through to 2008. Bell et al. (2022, Figure 5) report earnings volatility estimates for the period 1975–2020 derived from ASHE data, but these are not comparable with the other estimates cited because Bell et al. use an idiosyncratic measure of volatility (the absolute difference between the 90th and 10th percentiles of the distributions of one-year earnings changes).

We have produced up to date estimates of labour market volatility and earnings volatility, deriving series for $Mn(|a_t|)$ as well as $SD(a_t)$, and adding an ASHE-based series. See Figure 1 which confirms that earnings volatility remained constant through to 2021, according to both $SD(a_t)$ and $Mn(|a_t|)$ and all data sources. (The uptick in the $Mn(|a_t|)$ in 2022 reflects a sharp rise in inflation, which does not affect $SD(a_t)$.) Labour market volatility continued its secular decline, again according to both measures and both sources, though showing a Covid-19-related uptick after 2019. One factor behind the decline is a fall in the rate of UK labour market ‘churn’: quarterly inflows to and outflows from employment as a proportion of total employment have been declining over the past 25 years. See Brewer, Cominetti, and Jenkins, (2025, Figure 4).

<Figure 1 near here>

The estimates shown in Figure 1 are annual averages for the quarterly LFS and for the BHPS/UKHLS: small sample sizes mean that estimates broken down by month or quarter of interview exhibit substantial variability (though the general trends are the same as shown). One of the advantages of the administrative data to which we now turn is that sample sizes are very large – sampling variability is a negligible issue.

5.2. PAYE data: HMRC and Linked-ASHE files

Our research is based on the data that UK employers are mandated to report to the tax authorities (His Majesty’s Revenue and Customs, HMRC) each time they pay their employees using the Pay As You Earn (PAYE) Real Time Information (RTI) system. PAYE refers to employers’ withholding from employee pay packets the employee’s liabilities for income taxation and National Insurance contributions (NICs), with employers also passing the monies withheld on to HMRC. RTI means what it says and is based on regular electronic filing, replacing, since April 2014, a paper-based system with end of financial year reconciliations.

The earnings data are for a 1% random sample of the employees in the PAYE system who are in the target sample for ASHE. The selection of employees each year for ASHE is random, based on the (same) last two digits of an individual’s National Insurance number,

analogous to a Social Security number in the USA. For more information about ASHE, see e.g. Bell et al. (2022) and references therein.

The earnings information is derived from all the payslips submitted to HMRC between 6 April 2014 and 5 April 2019, i.e., covering the five financial years (FYs) prior to the onset of Covid-19, 2014/15 through 2018/19. For individuals with multiple jobs, we combine payslip information for a given pay period to derive a measure of total gross (pre-tax) employee earnings from all jobs ('pay') for that period. As in previous research, e.g., Blundell et al. (2024) for the USA, we interpret the absence of a payslip for an individual to mean 'not in employment' and set pay equal to zero.

We restrict analysis to individuals aged 20–59 to exclude young school-leavers and workers close to retirement. We also drop a small number of individuals whose records of within-year earnings are irreconcilable with an internal check variable cumulating earnings. For these cases, we dropped the data for all months within the relevant FY from the analysis samples.

We do not have access to all the information on a payslip. Pay received by an employee may include basic pay, overtime, bonuses and commissions, backdated pay increases, tips passed on by the employer, or holiday pay. We only have the pay total and do not have work hours for hourly paid workers. Although we do not have as much information about pay composition as Ganong et al. (2024), we do have data about workers linked in from ASHE, as described shortly.

Because employees are paid at different frequencies, we need to choose a period over which to calculate earnings to examine volatility among all employees. We rule out aggregation to the annual level because we wish to look at sub-annual earnings changes. In terms of how the worlds of employment, cash transfers from the government, and household bills, work in the UK, a 'monthly' reference period is the natural option. Moreover, the most common pay frequency in our main analysis sample is monthly (or bi-monthly), accounting for around 72% of all person-month observations in our main analysis sample. Hence, it is for the employees with a weekly pay period type that our consistency adjustments are focused on. In short, we 'daily-ise' the earnings of people paid weekly and then aggregate their earnings over the days of each calendar month (so someone paid the same each week could be paid more in January than February). See Appendix A and WED's (2024) documentation for further discussion of our monthly pay variable derivations.

We convert nominal monthly earnings amounts to real income terms (March 2019 prices) using the monthly Consumer Prices Index including owner-occupiers' housing costs

(CPIH). All the studies of earnings volatility based on annual earnings have used real earnings amounts, so too did the ASHE-based GRID study for the UK by Bell et al. (2022). For monthly earnings changes it is arguably more appropriate to use nominal monthly earnings, which is what Tomlinson (2018) does, but we have five years of data rather than one. See Appendix A for more discussion, and evidence that real and nominal changes provide similar distributions of monthly earnings changes. This was a period of particularly low inflation, with the monthly increase in the CPIH between April 2014 and March 2019 about 0.12% on average.

We do not pre-adjust earnings for differences by age. Within-year earnings variations with age are negligible, and many earlier studies (using an annual perspective) have reported that raw and adjusted estimates led to the same conclusions.

We work with two analysis datasets. The *HMRC file* contains the PAYE pay data and some information about employee characteristics, such as age, sex, payslip type, and whether has more than one job. The WED team determined age and sex by links to longitudinal ASHE data for survey years 1997 through 2022, so they can be derived even if a person with RTI data did not appear in the ASHE surveys for 2015–2019 that we use. Age is missing for a very small number of cases, i.e., no ASHE data were found for them despite their being in the target sample.

The *Linked-ASHE file* is our second analysis file. It contains PAYE pay data plus information from ASHE about the characteristics of employees and their jobs that has been linked to individuals in the HMRC file. ASHE collects data from employers about the characteristics of employees and their jobs, with most measures referring to the April of the ASHE survey year, from which a panel dataset can be created. ASHE variables include characteristics such as occupation and industry, whether hourly paid, or working on a temporary contract, and whether in the public/private sector. We spread the characteristics information from ASHE forward from April over the rest of the financial year, acknowledging that some characteristics may change within the year, and we cannot observe these changes. At the same time, we would emphasise the advantages of having characteristics data to provide informative breakdowns of the volatility patterns we observe.

A limitation of the Linked-ASHE data is that the linkage rate to HMRC data is well short of 100% and successful linkage is non-random. According to WED’s documentation, the low match rate is mostly due to ASHE only achieving around two-thirds of its sampling target (2024, p. 7). This may be because of employer non-response per se. (The WED team’s revised ASHE survey weights aim to adjust for this; we use these weights for all analysis of

Linked-ASHE file data.) Another important issue is that individuals with low employment attachment are less likely to be linked because they must be working in the April of a FY for an employer to be liable to file an ASHE response about them.

Our comparisons of the HMRC and Linked-ASHE files indicate that the latter leads to lower volatility estimates. For example, for our analysis samples and pooling the data for all 60 months, the HMRC file estimate of $Mn(|a_t|)$ is 0.150 for those working in months $t-1$ and t , and the labour market volatility $Mn(|a_t|)$ is 0.184, whereas the corresponding estimates from the Linked-ASHE file are 0.136 and 0.164. Looking at the full distribution of arc percentage changes among those working in months $t-1$ and t , the main differences between the HMRC and Linked-ASHE distributions are in the prevalence of small versus large earnings changes. For example, pooling all 60 months' data, 47.5 per cent of arc percentage changes are between -0.01 and $+0.01$ in the HMRC file but 36.6 per cent in the Linked-ASHE file. Looking at large changes, 8.5% of changes are greater than 0.25 and 8.2 per cent are less than -0.25 in the HMRC file, whereas the corresponding fractions in the Linked-ASHE file are 8.3 per cent and 10.9 per cent. See Appendix Figure E1 for more details.

The lower volatility estimates from the Linked-ASHE file arise in part because workers with lower attachment to employment are less likely to be successfully linked. For example, in the HMRC file, those working in months $t-1$ and t contribute 76.5 per cent of the person-month observations, whereas in the Linked-ASHE file the corresponding fraction is 92.9 per cent. Earnings volatility is also lower according to the Linked-ASHE file because it contains fewer observations from the poorest tenths of the HMRC file earnings distribution (in month $t-1$) and earnings volatility is relatively high for the poorest tenth (as we show below). Reassuringly, however, we find broadly similar volatility differentials across subgroups defined using characteristics present in both the HMRC and Linked-ASHE files (age, sex, number of jobs, pay frequency type), and so we expect that volatility breakdowns by characteristics only present in the Linked-ASHE file should also be reliable. See Appendix Table E1 and Figures E3–E9 for detailed cross-file comparisons.

After sample selections, the HMRC file contains data for an unbalanced panel of around 294,500 employees, of whom 74 per cent are present for all 60 months. (Present refers to observations with zero or positive employment earnings.) There are about 15.9 million person-month observations of which 12.1 million refer to individuals working in the current and previous month. The Linked-ASHE file, also an unbalanced panel, contains around 210,000 workers, 17 per cent of whom are present for all 60 months. It contains around 7.1 million person-month observations of which about 6.6 million refer to individuals working in

the current and previous month. The smaller fraction in the Linked-ASHE file of individuals employed all 60 months is another manifestation of its under-representation of those with low employment attachment by comparison with the HMRC file. Focusing our breakdowns by characteristics on individuals working at $t-1$ and t when analysing monthly volatility, and individuals working 12 months in a FY when analysing multiperiod volatility, mitigates this problem. This focus on working individuals also corresponds with the focus of most recent research on earnings volatility.

5.3 Pay volatility according to the PAYE data: a first look

We begin by presenting time series for monthly labour market and earnings volatility estimates over the five-year period, and then document differences across groups in more detail. Figure 2 contains three pairs of volatility series. Within each pair, the dashed line shows volatility, $Mn(|a_t|)$, calculated for all workers, including those with zero earnings at $t-1$ or t (labour market volatility), whereas the solid line shows volatility calculated only for workers with positive earnings at $t-1$ and t (earnings volatility). The vertical dashed lines in Figure 1 and subsequent charts identify April and November in each year, i.e., the months in which volatility has a peak and a trough.

The red series at the top of Figure 2 referring to monthly earnings changes over a one-year gap is the most comparable with earlier panel survey-based estimates for the UK, with the main difference being that we can use a moving one-year window whereas earlier research has reported only one estimate per year. The difference matters: with the rolling window a clear seasonal pattern emerges, with volatility increasing towards the end of the FY, peaking in March, and then declining again. Peaks and troughs aside, Figure 1 shows a small decline in volatility over the five-year period. The gap between the solid and dashed red lines is about 3 percentage points throughout the period, indicating that employment entries and exits account for about 7% of labour market volatility. The estimates are nearly twice as large as the estimates of $Mn(|a_t|)$ derived from UK panel surveys (see above). We attribute this difference to substantially better coverage of the top and bottom of the earnings distribution in our PAYE data.

<Figure 2 near here>

The blue series at the bottom of Figure 2 shows volatility using one-year changes in annual earnings (the 12-month sum over each FY). These are the first UK estimates of this kind – no UK household survey collects annual earnings data. Volatility measured this way is

a tiny fraction of volatility according to the other two measures – about one-twentieth of the ‘monthly earnings, a year apart’ series – and shows no trend.

The middle pair of series shows monthly earnings instability within each year. Clearly, there is substantial volatility in monthly earnings that individuals experience between one month and the next, and any apparent trend over the five-year period disappears. The estimates are around half the size of the corresponding ‘year apart’ series at the top of the chart reflecting the fact that lengthening the window over which change is calculated increases the chances of perceptible changes occurring. (This has long been known from the US PSID-based research comparing estimates based on one-year gaps and two-year gaps.)

In each of the five FYs, there is a similar pattern: volatility is lowest in November, reaches a peak in April (the first month of a FY), falls sharply in the following month, and then falls relatively gradually over the summer months to reach a trough in November. Volatility is almost as high in March as in April, but distinctly lower in May than April: each sawtooth is wider on its left than its right.

The amplitude of the within-year fluctuations is relatively large, with a minimum value of $Mn(|a_t|)$ for the ‘working at $t-1$ and t ’ series around 0.13 and the maximum around 0.18, compared to the period average of 0.15. Although the mean and median absolute change over the period is zero, around one-tenth of the arc percentage changes refers to a fall of at least 21 arc per cent and around one-tenth refers to a rise of at least 22 arc per cent. Recall too the relatively high prevalence of changes of at least 25 arc per cent. The distribution of changes is symmetric, but its density has much narrower shoulders than does the bell-shaped density of a normal distribution as also reported by the GRID project (Guvonen et al., 2022). See Appendix A for further details of the distributions of earnings changes.

All in all, although many workers experience little change in their earnings from one month to the next, there is a significant minority that experience large falls or rises, and there is also notable seasonality in changes.

The labour market and earnings volatility series move in parallel in all three pairs of series in Figure 2. Movements into and out of employment make a non-trivial contribution to labour market volatility: for the ‘monthly, a month apart’ series, these transitions account for around 25% of overall labour market volatility (which averages around 0.2). But the fraction has not changed over time.

The seasonality patterns and lack of secular trend shown in Figure 2 are robust to the choice of volatility index (complementing what we saw in Figure 1). In Appendix Figure B1

we show that, for month-to-month earnings changes, the mean absolute arc percentage change, the standard deviation of arc percentage changes, and the standard deviation of the change in log earnings, provide similar pictures about within-year seasonality and trends over the period. All three indices show a volatility peak in the April (with March volatility almost as large), a trough in the November of each year, and there is no volatility trend over the five-year period. The main difference between the series is that the SD-based measures show somewhat greater within-year fluctuations than does the mean-based one and, as expected, this is more apparent for labour market volatility than earnings volatility. The series amplitude is greater relative to the series average and the SD(log changes) index also shows a mini peak over the summer.

5.4 Monthly pay volatility: breakdowns by sex, payslip type, and earnings level

Figure 3 compares monthly pay volatility between men and women in terms of both labour market volatility (top series) and earnings volatility (bottom series). Seasonality appears more prominent than in Figure 2, but this is simply because the vertical axis scale has changed.

<Figure 3 near here>

Men's and women's volatilities are similar in terms of levels and seasonality. For both sexes, volatility peaks in April and troughs in November. The most perceptible differences are that men have a more distinct decline in volatility between May and November than women, whose series displays a mini peak in September, and these differences are most apparent in the labour market volatility series. In other words, transitions into or out of employment are more pronounced for women than men at that time of year. Our conjecture is that these differences are due to mothers being more likely than fathers to not work or work fewer hours during the summer school holidays.

The breakdowns that follow are restricted to individuals with positive earnings at months $t-1$ and t because the labour market and earnings volatility series move in step through the five years, and the numbers with employment entries and exits is small. The proportion of person-month observations with an arc percentage change of 2 (employment entries) is only 1.7% and the proportion with -2 (employment exits) is also 1.7% (Appendix Figure A1).

Figure 4 shows how volatility patterns differ according to payslip type. Monthly payslip types are by far the most common (around 73% of all person-month observations) with weekly payslip types comprising almost all the rest (23%), and their volatility exhibits different seasonality patterns. For the monthly payslip group, $Mn(|a_t|)$ is broadly constant at

around 0.15 across the period, except for jumps of at least 5 percentage points (ppt) in March and April of each year. In contrast, for the weekly payslip group, $Mn(|a_t|)$ is about two-thirds larger on average, fluctuating between 0.25 and 0.30. There are peaks in March and April (as for the monthly-paid group), but troughs in August and November are much more discernible. For the weekly payslip group, March 2019 also stands out because of the very large rise in $Mn(a_t)$, explained by the fact that this was a month with five pay days for weekly-paid workers, assuming they were paid on a Friday (the situation for around 80% of employees; see Murphy 2021, p. 11).

<Figure 4 near here>

Figure 5 documents how earnings volatility levels and seasonality differ across tenths of the distribution of earnings in base month $t-1$. To reduce chart clutter, series for the middle four tenths are not shown, but the action is at the very top and very bottom in any case. (D1 refers to the bottom tenth, D10 to the top tenth.) The chart makes clear that the sawtooth seasonal pattern seen in aggregate (Figure 2) is driven by earnings changes for the richest tenth of earners, almost all of whom would have a monthly payslip type (Figure 4). For this group, the volatility series has a very large amplitude by comparison with those for the other groups – nearly 20 ppt – with a spiky peak in March and April and trough in September, October, and November. The bottom earnings tenth is distinctive because it has high volatility levels on average (substantially more than even the second tenth), systematic seasonal variation is less perceptible, and the amplitude of the fluctuations is relatively small. As one moves up the earnings distribution from the poorest tenth, volatility levels fall sharply until reaching the top tenth which has its own distinctive seasonal volatility pattern.

<Figure 5 near here>

Figure 6 provides complementary details about the differences in patterns observed in Figure 5. Panel (a) shows the fractions of individuals in the poorest tenth of earnings in month $t-1$ with a large increase or a large decrease in pay from month $t-1$ to t , where ‘large’ means ‘at least 25%’. Panel (b) shows the corresponding series for the richest tenth at $t-1$. For the bottom decile group, the fraction with a large increase in any given month is greater than the fraction with a large decrease, and the reverse is true for the top decile group. This is partly a reflection of regression to the mean: on average, earnings for those at the bottom (top) are likely to go up (down). But there is more going on; the pairs of series also show other interesting differences.

For the bottom tenth, the amplitudes of both series are relatively small, about 5 ppt, with two exceptional months. January and especially February are when the fraction with a

relatively large decrease is greatest and the fraction with a large increase is smallest. The February feature might be due in part to the way our monthly earnings variables are created (see Appendix A about the impact of February having fewer days than January).

<Figure 6 near here>

In contrast, for the top tenth of earners, March and April stand out. The fraction with a large rise is between 5% and 8%, except in March when it jumps to over 15%. The fraction with a large fall is between around 12% and 20%, except in April when there is an exceptionally large spike peaking at around 33%. More generally, the group's large decrease and large increase series move in tandem and more obviously so than for the bottom tenth: the temporal profile of the large decrease series is essentially the same as the profile for the large increase series except that it is shifted one month to the right. For example, look at turning point pairs in June and July, and December and January, in addition to March and April.

An explanation for the large March rise – large April fall pattern is the seasonal nature of bonus payments in the UK. According to the ONS, “[t]ypically, the majority of large bonuses are paid between December to March each year, mainly, but not exclusively, in the financial and insurance activities industry” (2017, p. 7). As it happens, the bonus season peaks in March not only for this industry but also the rest of the economy (ONS, 2017, Figures 5a, 5b). And large bonuses are more likely to be paid to workers in the top tenth of earnings. Figure 7 shows the seasonal pattern to bonus payments across our five-year sample period. The chart shows that bonuses are particularly important in the financial and business services and yet there is also a marked seasonal pattern across all sectors in the economy. For example, there are some low-paying sectors for which a similarly shaped (but less extreme) seasonal pattern is also observed, notably ‘wholesale, retail, hotels and restaurants’.

<Figure 7 near here>

Although the exact size of large bonuses is likely to be uncertain to their recipient, their occurrence and timing is relatively predictable and so too is the fall in total pay in the following month. Bonus-driven monthly volatility for an affluent group is unlikely to represent income risk of social concern. Of greater concern is the high level of volatility experienced by those at the bottom of the pay distribution, and for them the seasonal pattern is much less systematic. We now consider which other characteristics identify groups with high volatility levels, showing that many are commonly associated with being low paid.

5.4. Monthly pay volatility: additional subgroup breakdowns

For brevity and noting the lack of trend over the sample period, our additional subgroup breakdowns are based on month-pair observations pooled for all 60 months, though we also comment on distinctive subgroup seasonal patterns to volatility drawing on the full month-by-month subgroup charts in Appendix B. We provide breakdowns by sex, age, selected low- and high-paid occupations, and various job characteristics such as whether in the public sector, whether the contract is permanent or temporary, and hourly paid or not, whether the worker has more than one job, and whether has been in the job for more than a year. Almost all characteristics are derived from ASHE.

Table 1 shows volatility estimates using three measures to check the robustness of subgroup volatility differentials to summary metric: $Mn(|a_t|)$, $SD(a_t)$, and $Var(a_t)$. We also report subgroup estimates expressed as a percentage of the corresponding estimate for all observations (top row) to more easily identify relatively large differences from the average. Also shown are the relative sizes of the subgroups (number in subgroup expressed as a percentage of the total sample size). Recall that $Mn(|a_t|)$ for all persons equals the size-weighted sum of the subgroup $Mn(|a_t|)$ estimates. Thus, a subgroup can make a relatively large contribution to aggregate volatility if it has relatively large volatility or has moderately large volatility but is relatively large. For $Var(a_t)$, the decomposition only holds approximately.

<Table 1 near here>

The first breakdowns in Table 1 refer to sex and age. Clearly, there are few differences by sex (as Figure 3 would lead us to expect), but volatility decreases with age up until around age 45 after which it is much the same. In particular, workers aged 20–24 have very high volatility compared to all other groups. For example, their $Mn(|a_t|)$ estimate is around 27% higher than the estimate for workers aged 25–29, and the latter group’s estimate is about 7% higher than the estimate for those aged 30–34. Differences in volatility across older groups are much smaller. (The other two measures tell the same story about subgroup volatility rankings but using $Var(a_t)$ magnifies the numerical differences.) Appendix Figure B2 shows that the 20–24 group have a distinctive seasonality profile. As for the older groups, volatility is lowest around November but, differently, March and April are not the peak volatility months. There is a mini peak in April and a small fall in May, but volatility then rises to reach a peak in September and October. It is unclear what is driving this, but one potential explanation is that it reflects movements from full-time work to part-time work at

the end of a summer, perhaps also associated with taking up further or higher education or training.

The next panel summarizes the differentials by earnings decile group already pictured in Figure 5. Seasonality differences aside, volatility is much greater in the top tenth of the earnings distribution than the pooled average – some 40 per cent larger according to $Mn(a_t)$ – but not as large as volatility among the poorest tenth, who have a volatility level more than twice the pooled average. Correspondingly, volatility is well below average for decile groups 4 to 9. But what about volatility differences within low and high paid groups?

Table 1 summarises earnings volatility differences across low-paid occupations. We consider the 14 defined thus by the UK Low Pay Commission (2017 definitions based on SOC2010 occupational classifications) which together account for around 31% of our analysis sample observations. Retail is by far the largest low-paid occupation (8.3% of the sample), with the next largest low-paid occupation, hospitality, around half the size (4.7%). Cleaning, social care, and storage, are the next largest groups (between 2.6% and 2.9%). The largest estimates of $Mn(|a_t|)$ are for social care and hospitality, almost 50% and 45% larger than the aggregate estimate, though hospitality accounts for a greater share of overall volatility than social care because of its larger size. However, by far the largest share of total volatility is accounted for by Retail because its relatively large $Mn(|a_t|)$ estimate is weighted by a much larger size. In contrast, volatility is also high for Leisure but the sample fraction in this occupation is tiny.

If the volatility measure is changed to $SD(a_t)$, Table 1 shows that the volatility ranking of low-paid occupations is broadly the same but with some differences. For example, social care and hospitality swap places in the volatility ranking. Textiles has volatility above the aggregate estimate according to $Mn(|a_t|)$ but below it according $SD(a_t)$ and $Var(a_t)$. The latter two measures rank the subgroups identically of course but yield different subgroup relativities because they use different scales. $Var(a_t)$ provides subgroup volatility relativities that are on roughly same scale as $Mn(|a_t|)$, whereas those for $SD(a_t)$ are more compressed. (This is a general feature of many but not all breakdowns in Table 1.)

Across low-paid occupations there is substantial heterogeneity not only in volatility levels but also in seasonality patterns. Appendix Figure B3 shows, for example, that the seasonal pattern for Retail is like the overall average pattern, but the amplitude of the fluctuations is larger. The shape of the within-year profile differs from the aggregate one for several low-paid occupations, including hospitality and childcare (with peaks after December and summer months).

The next panel of Table 1 refers to selected high-paid occupations. These account for around 15% of our overall sample and so a larger group than the top decile group considered earlier (which perhaps explains why not all groups have above-average volatility). Managers and directors working in financial institutions have by far the greatest volatility on average, around 30% larger than the overall $Mn(|a_t|)$, and they have a different seasonal volatility pattern. There is a March-April spike around four times larger than average (as the earlier charts lead us to expect), there is a smaller peak in June and July, and no distinct trough in November (Appendix Figure B4). In contrast, the other five high-paid groups have volatility patterns that mimic the overall pattern in terms of peak and trough months. However, the financial institution managers and directors are a tiny group and so make little contribution to aggregate volatility. In contrast, the sales, marketing, and related occupations group is around 16 times larger, and makes a much greater contribution to overall volatility. So too, do functional managers and directors, because their above-average volatility is coupled with a larger subgroup size.

Differences between employees on monthly and weekly payslip types reflect those already discussed (Figure 4). Volatility is much greater on average for workers with multiple jobs, or working on a temporary rather than permanent contract, with $Mn(|a_t|)$ around 70% larger than the overall aggregate value. These differentials are greater than for any of the low-paid occupation subgroups and the groups are also relatively large: the multiple job and temporary contract worker groups each form between 6% and 7% of the sample. Moreover, unlike many other groups, these two experience a volatility peak during the summer: see Appendix Figures B7 and B8.

Being hourly paid is associated with higher volatility and the group is relatively large, almost 40% of the sample. In contrast with those not hourly-paid – mostly salaried workers – pay can vary with changes in weekly work hours and working at overtime pay rates. As a result, according to $Mn(|a_t|)$ volatility is around 40% larger for hourly-paid workers compared to those who are not. Working part-time is also associated with higher volatility, around 38% higher than for full-time workers. Employees remaining in the same job for at least a year experience earnings volatility about a quarter lower than those who change jobs.

Table 1 also shows that there is a volatility premium to working in the public sector (nearly one-fifth of the sample). Its $Mn(|a_t|)$ estimate is around one third lower than private sector workers (whose volatility is just above the sample average). Appendix Figure B6 shows that within-year variation in volatility is essentially a private sector phenomenon. It is

for that group that we see the distinct sawtooth shape shown for all workers in Figures 2 and 3.

5.6. Multiperiod volatility measures

We have also derived estimates of volatility using monthly earnings over a 12-month period, a multiperiod approach in contrast to the period-pair approach relied on so far. (See Brewer, Cominetti, and Jenkins, 2025, for more extensive analysis.) Here we are interested to see whether the longer assessment period makes a substantive difference to which groups are identified as having the greatest (or least) volatility. For brevity, our estimates and commentary are presented in Appendix C. The main conclusions are that subgroup volatility rankings are essentially the same as for monthly volatility (Table 1), but the magnitudes of the differentials depend on whether the multiperiod volatility measure accounts for all deviations from longitudinal average earnings (average CV measure) or focuses on the prevalence of multiple large deviations (4+ monthly deviations at least 25% above or below longitudinal average earnings). Differentials are larger for the latter measure.

5.7. Factors associated with high levels of pay volatility

Our analysis of HMRC data shows that month-to-month pay volatility is relatively high in the UK, perhaps of a similar order to that in the USA, and there is a distinct within-year seasonal pattern with volatility peaking in March and April and at a minimum in November. Earnings volatility is greatest among those at the very top or the very bottom of the earnings distribution. Other factors associated with high pay volatility levels include working in a low-paid occupation, in a private sector job, paid by the hour and being paid weekly, on a temporary contract, and being aged 20–24. The multiperiod volatility breakdowns point to the same factors as well.

These associations raise several questions. One is whether the factors have independent effects, e.g., what is the association between volatility and low-paid occupation if one adjusts for type of employment contract or whether hourly paid? The main conclusion from our regression analysis presented in Appendix D is that each of the characteristics associated with higher volatility in our univariate analysis remain important when adjusting for other characteristics.

A second and related question is how volatility profiles look for individuals with specific combinations of characteristics. In Appendix Figure D1, we contrast the volatility experiences of five ‘types’ with those of all other individuals, where types are constructed

using assumptions about age, whether hourly paid, payslip type, and occupation and sector. Amongst other things, the chart shows the volatility profile for someone aged 45–49, with monthly payslips and not hourly paid, who works as a higher or further education professional, in other words the profile for someone in a relatively secure job. For those working at $t-1$ and t , we estimate $Mn(|a_t|)$ to be around 0.05 (60 month pooled average), which is substantially lower than the sample average and seasonality is low by comparison with All workers. This provides a useful lower bound benchmark for what counts as ‘low’ volatility.

A third question concerns the drivers of high volatility, especially how it relates to variations in pay rates versus hours variations (we cannot observe the distinction with our PAYE data). For earners in the top tenth, we have argued that bonuses play an important role in explaining that group’s high volatility. But what about other earners, especially those in the bottom tenth? We have also shown that earnings volatility is higher for workers who are hourly paid, have a weekly payslip type, etc. For all these groups, the most obvious explanation of their high volatility is instability of work hours.

It is hard to draw conclusions about changes in the instability of work hours over time in the UK because of measurement challenges. One can look at the fraction of employees paid at a fixed hourly rate who say their hours vary. This is relatively high – around 18% in 2014 – though the fraction declined subsequently, to around 16% in 2018 (Brewer et al., 2025, Figure 5). Employees on zero-hours contracts (ZHCs) face more insecurity and work hours may fall to zero: their employers do not guarantee any minimum number of working hours and only offer work when it is available. There was a marked increase in the fraction of employees on ZHCs over the 2010s, going from around 2% in 2014 to around 3% in 2018: see Brewer et al. (2025, Figure 5). Farina, Green, and McVicar (2019) argue this rise was this was partly linked to LFS respondents being more aware of zero-hours contracts from around 2012/13, but they also state that the LFS likely underestimate the number of workers whose main job is a ZHC, especially if other types of no-guaranteed-hours-contract jobs are counted.

We cannot identify employees with a ZHC in our PAYE data. However, it is notable that the industries in which workers experience the largest earnings volatility are also those which make the greatest use of zero-hours contracts. This is illustrated by Figure 8 which plots $Mn(a_t)$ against ZHC prevalence by industry and shows a clear positive relationship. Observe too that “ZHCs have become increasingly concentrated at younger ages, among full-time students, among migrants, among black and minority ethnic workers, in the private

sector, in personal service and elementary occupations, and in the distribution, accommodation and restaurant sector” (Farina et al., 2019, p. 509). Most of the groups mentioned are ones that we showed earlier to have high volatility levels. (We exclude students and do not observe migrant or ethnic minority status.)

<Figure 8 near here>

6. Policy responses, earnings volatility, and the UK’s Universal Credit

We motivated our analysis of earnings volatility by arguing that it provides a signal of income risk. The standard caveat by economists regarding the interpretation of empirical estimates such as ours is stated by Shin and Solon, who wrote that:

[I]dentifying the risk associated with earnings changes will require further information on whether the changes were or were not anticipated (or even purposively chosen) and whether the affected individuals were or were not insured against the changes (through such means as transfer programs, saving/borrowing, or family labor supply adjustments). Assessing the welfare implications of changes in measured earnings volatility ultimately will require answers to these difficult questions. (Shin & Solon, 2011, p. 977.)

Their conclusion is undeniable, but it should not be interpreted as suggesting policy responses to counter volatility’s downsides should be put on hold. We remarked earlier that it is the relatively high levels of monthly volatility for people at the bottom of the earnings distribution that are likely of greater social concern than similarly high volatility experienced by high-paid workers receiving bonuses. There is sufficient evidence about the nature of earnings (and income) volatility to justify policy responses to help low-income families better cope with uncertain income flows.

Moreover, low-income families commonly express a desire for less income instability, disliking the uncertainty they report. For the UK, see e.g. the qualitative responses reported by Angsten Clark et al. (2024), Biosca et al. (2020), Griffiths and Wood (2024), Hills et al. (2006), Miller et al. (1989), and Tomlinson (2018). Moreover, the lowest-paid fifth of employees in 2017 were nearly four times as likely to say they were ‘very anxious’ about ‘unexpected changes to my hours of work’ as the highest-paid fifth (11 per cent versus 3 per cent): see Brewer, Cominetti, and Jenkins (2025, Figure 6). For the USA, a Pew Trust

(2015) survey asking whether respondents would prefer to have financial stability or to move up the income ladder, 92% chose security, an increase of 7 percentage points since 2011.

One of the roles that welfare states may play is to provide in-work transfers to cushion the effects of earnings volatility, especially large earnings decreases, and to do so in a timely fashion. The problem is that the cash transfer system supporting low-income families has features that may also exacerbate earnings volatility problems. We elaborate these issues in the rest of this section, reiterating arguments about the UK made by Tomlinson (2018). For further elaboration, see Johnson et al. (2025) and especially the detailed qualitative research by Griffiths and Wood (2024).

Universal Credit (UC) is the UK's principal form of means-tested state cash support for working-age families whether in or out of work, progressively rolled out since 2013 and so covering our data period. To underline the importance of UC support for working families, observe that nearly 40% of the 6.4 million UC claimants in December 2023 were in paid work or living with a partner who was working (Griffiths and Wood, 2024, p. vi).

Importantly, the UC assessment period is one month, with monthly reckoning up combined with potential clawbacks of past overpayments, where means are assessed using data from HMRC's PAYE RTI system described earlier and other factors such as household composition. In contrast, the former Tax Credit system made the same payment to an eligible family regularly throughout the year, but there was an annual reckoning up with possible clawbacks of overpayments. Compared to the Tax Credits, UC's shorter reference period, one month rather than one year, allows potentially greater responsiveness to changes in circumstances, as illustrated during Covid-19 when additional government support was provided to UC recipients very quickly.

There are UC reference period issues that apply specifically to workers paid weekly or some weekly type. As shown earlier, this group is more likely to experience higher volatility than monthly-paid workers and are more likely to be UC recipients because they are lower paid on average. Workers paid each week have four months in the year in which they receive five pay packets, those paid fortnightly have two months in the year in which they receive three pay packets, and those paid four-weekly have one month in the year with two pay packets. Months with unusually high monthly pay will lead to downward adjustments in the next UC payment even if the same amount is received in each weekly pay packet, and the problem is compounded if earnings are volatile from weekly pay packet to pay packet. This process raises the chances that low-income families will find it hard to cope, especially if the

UC adjustment means loss of all UC entitlements in which case a new application must be made to receive any payment, and there is a five-week wait period until the first payment.

Monthly-paid workers have UC reference period issues too because UC's monthly assessment period is not aligned with the month corresponding to pay dates. On the one hand, UC assessment period dates are determined by the day on which claimants first apply for UC. On the other hand, the last day of the month is the most common payday for monthly payrolls (~ 40%), and Fridays most common for most other pay frequencies (Murphy, 2021). To take an example from Tomlinson (2018) that remains relevant, suppose the UC assessment period starts on the 18th of the month, UC is paid on the 25th of the month, and pay is received on the last day of month. Hence, e.g., a UC payment on 25 February is based on January's pay, which was most likely received on the 31st of January (or the last Friday of the month), i.e., with significant time lag. Even more importantly, it is possible to show that fluctuations in pay can be amplified by UC, i.e., the volatility of income from pay and UC combined can be greater than the volatility of pay alone.

In sum, gains from speedier responsiveness to circumstances can be offset by the misalignment of reference periods across income sources. UC payment levels are also important because impoverished individuals and families are less able to manage uncertainty through budgeting and planning, including borrowing and saving. There also issues related to deductions from UC payments for loan repayments, benefit overpayments, and third-party debts. Although UC may be potentially helpful from a financial management perspective, problems remain if insufficient financial support remains to cover everyday living costs. All in all, there is a good case for revising UC's rules regarding reference periods to take account of the substantial earnings volatility faced by many low-income working families.

Measures directed at employers would also help. Their actions also matter because of their key role in the PAYE process, frequently and regularly reporting pay to HMRC, and including UC payments in the pay packets of UC-eligible workers. There can be unintended adverse effects on worker's UC eligibility if employers provide back-dated pay increases or performance bonuses, file their PAYE returns late, or make corrections to initial earnings reports to HMRC.

These arguments refer specifically to the UK context, but there are similar issues for other countries about how the administration of cash transfer systems could be reformed to mitigate the impacts of undesirable income instability. For the USA, Morduch and Zwicky argue that "the inclusion of most transfers ends up increasing, as opposed to decreasing, the standard deviation of monthly income, which is a sign that transfers exacerbate rather than

buffer the income volatility experienced by US Financial Diaries households” (2017, p. 412). Romich and Hill’s (2017) review points to volatility related issues at three stages in the administration of transfers – initial application, within receipt periods, and (re)certification.

7. Summary and conclusions

We have reviewed research about income and earnings volatility and provided new UK evidence about the latter using high-quality administrative record data. Our review argues that, although household income volatility is of greater interest than the earnings volatility from a social welfare point of view, there are fewer up-to-date national or cross-national studies of income volatility than for earnings. Whether considering income or earnings volatility, the USA stands out as a high volatility country relative to the UK and other high-income countries and, although the large increase in US volatility over the 1970s and 1980s did not continue, it has not fallen either. Volatility remained fairly constant in other countries as well, according to the most recently available data.

Almost all volatility research to date has considered volatility from an annual perspective, and yet many low-income families must manage regular monthly expenditure commitments with incomes that are unstable from one month to the next or even from one week to another (as corroborated by the qualitative research cited earlier). In this paper, we have provided new evidence about within-year earnings volatility from one month to the next. We have shown there are marked seasonal patterns to volatility in aggregate. Earnings volatility is highest for employees in the highest tenth or the lowest tenth of earners but otherwise with different features. For those at the top, high volatility and its seasonal pattern likely reflects the payment of bonuses. For those at the bottom (and overlapping groups such as young people, individuals working in a low-paid occupation, etc.), the high volatility is more generic and has a different and more muted seasonality. It likely reflects instability of work hours, including that arising from zero-hours contracts.

These findings about low earners have relevance to the design of cash transfer support in the UK because the monthly reference periods that UC uses do not align with many earners’ pay periods and so can exacerbate volatility problems. Researchers have cited analogous issues for the USA. What the situation is for other countries is an open question to be answered by future research.

Volatility also has other costs that impact low-income families in particular, including psychological and emotional costs and reduced opportunities to make longer-term financial plans (Angsten Clark and Otulana, 2024). With distributionally-skewed volatility, we continue to have what Morduch and Schneider referred to as “hidden inequality – an inequality in exposure to risk and in access to dependable ways to cope” (2017, p. 13).

Our paper also speaks to the strengths and limitations of different types of longitudinal data. A big advantage of household panel surveys is that they collect data about multiple household income sources (including labour earnings, capital income, and transfers) and the characteristics of household members, thereby enabling researchers to document household income volatility levels and trends and relate them to income components and household composition. However, few panel surveys enable examination of volatility within a year and there are issues of sample size, attrition, and measurement error.

Many administrative data sets address the latter three issues, and some contain observations at a higher frequency than annual. But administrative datasets derived from social insurance contribution records refer only to employee earnings and contain little information about worker characteristics, and those based on bank account transaction records or payroll processing data may not be representative, and it is more difficult to link records for individuals from the same family to look at income volatility. Linking administrative earnings records to survey respondents, as has been done with US CPS-ASEC and SIPP data is a valuable compromise, though implementing such linkages is not necessarily straightforward, dependent on country-specific laws about data protection and consent and ensuring that the linkages made are accurate.

Our HMRC PAYE data illustrate these various data issues. They provide very large nationally representative samples of employee earnings and earnings changes from payslip to payslip within a year, but they do not currently contain all the information reported by employers to HMRC so we cannot distinguish between volatility in base pay and volatility arising from variations in work hours, overtime pay rates, etc. Enhancing the data with this information, linking to other income data held by HMRC from tax returns (including self-employment data), and to other household members, while also extending the data series forward (and backwards) would substantially increase opportunities to document and understand earnings and income volatility, and many other issues. At the same time, as US experience shows, it would be valuable to link HMRC data from tax returns, as well as other government agencies’ data about cash transfers, to existing UK panel surveys such as the UKHLS and the LFS.

Improvements to datasets for specific countries may have the disadvantage of making cross-national comparisons using comparable data for multiple countries more difficult. The GRID project illustrates the potential for comparative research on earnings volatility using administrative data, albeit from an annual perspective. It would be great to see improvements in data infrastructure that would allow similar research on household income volatility, and also facilitate analysis of within-year volatilities of earnings and income.

8. References

- Acs, G., Loprest, P., and Nichols, A. (2009). ‘Risk and recovery. documenting the changing risks to family incomes’, Brief 8, Urban Institute, Washington DC.
<https://www.urban.org/sites/default/files/publication/32966/411890-risk-and-recovery-documenting-the-changing-risks-to-family-incomes.pdf>
- Angsten Clark, A. and Otulana, S. (2024). *Fluctuation Nation*, Real Accounts report, NEST Insight. <https://www.nestinsight.org.uk/wp-content/uploads/2024/11/Fluctuation-Nation-lifting-a-lid-on-the-millions-of-people-managing-a-volatile-income.pdf>
- Angsten Clark, A., Otulana, S., and Phillips, J. (2024). *Balancing Points: the struggle for stability*. Real Accounts Impact Brief #2, NEST Insight.
<https://www.nestinsight.org.uk/wp-content/uploads/2024/05/Real-Accounts-Impact-Brief-2-Balancing-Points.pdf>
- Avram, S., Brewer, M., Fisher, P., and Fumagalli, L. (2022). ‘Household earnings and income volatility in the UK, 2009–2017’, *Journal of Economic Inequality*, 20, 345–369.
<https://doi.org/10.1007/s10888-021-09517-3>
- Baker, M. and Solon, G. (2003). ‘Earnings dynamics and inequality among Canadian men, 1976–1992: evidence from longitudinal income tax records’, *Journal of Labor Economics*, 21, 289–321. <https://doi.org/10.1086/345559>
- Barnia, N. and Leete, L. (2009). ‘Monthly household income volatility in the U.S., 1991/92 vs. 2002/03’, *Economics Bulletin*, 29, 1–13.
<https://www.accessecon.com/Pubs/EB/2009/Volume29/EB-09-V29-I3-P59.pdf>
- Bartels, C. and Bönke, T. (2013). ‘Can households and welfare states mitigate rising earnings instability?’, *Review of Income and Wealth*, 59, 250–282.
<https://doi.org/10.1111/j.1475-4991.2012.00497.x>

- Beach, C. M., Finnie, R., and Gray, D. (2010). ‘Long-run inequality and short-run instability of men’s and women’s earnings in Canada’, *Review of Income and Wealth*, 56, 572–596. <https://doi.org/10.1111/j.1475-4991.2010.00406.x>
- Bell, B., Bloom, N., and Blundell, J. (2022). ‘Income dynamics in the United Kingdom and the impact of the Covid-19 recession’, *Quantitative Economics*, 13, 1849–1878. <https://doi.org/10.3982/QE1872>
- Biosca, O., McHugh, N., Ibrahim, F., Baker, R., Laxton, T., and Donaldson, C. (2020). ‘Walking a tightrope. using financial diaries to investigate day-to-day financial decisions and the social safety net of the financially excluded’, *The Annals of the American Academy of Political and Social Science*, 689, 46–64. <https://doi.org/10.1177/0002716220921154>
- Blundell, R., Pistaferri, L., and Preston, I. (2008). ‘Consumption inequality and partial insurance’, *American Economic Review*, 98, 1887–1921. <https://doi.org/10.1257/aer.98.5.1887>
- Blundell, R., Bollinger, C. R., Hokayem, C., and Ziliak, J. P. (2024). ‘Interpreting cohort profiles of lifecycle earnings volatility’. *Journal of Labor Economics*, forthcoming. <https://doi.org/10.1086/732667>.
- Bossert, W. and D’Ambrosio, C. (2011). ‘Measuring economic insecurity’, *International Economic Review*, 54, 1017–1030. <https://doi.org/10.1111/iere.12026>
- Brewer, M., Cominetti, N., and Jenkins, S. P. (2025). *Unstable Pay. Earnings Volatility in the UK*. Resolution Foundation report. <https://www.resolutionfoundation.org/publications/unstable-pay/>
- Cappellari, L. and Jenkins, S. P. (2014). ‘Earnings and labour market volatility in Britain, with a transatlantic comparison’, *Labour Economics*, 30, 201–211. <https://doi.org/10.1016/j.labeco.2014.03.012>
- Dahl, M., DeLeire, T., and Schwabish, J. (2011). ‘Estimates of year-to-year volatility in earnings and in household incomes from administrative, survey, and matched data’, *Journal of Human Resources*, 46, 750–774. <https://doi.org/10.3368/jhr.46.4.750>
- Debacker, J., Heim, B., Panousi, V., Ramnath, S., and Vidangos, I. (2013). ‘Rising inequality: transitory or persistent? New evidence from a panel of U.S. tax returns’, *Brookings Papers on Economic Activity*, 1, 67–142. <https://doi.org/10.1353/eca.2013.0004>

- Druehdahl, J., Graber, M., and Jørgensen, T. H. (2023). ‘High frequency income dynamics’, unpublished paper, University of Copenhagen. <https://michael-graber.github.io/pdf/Druehdahl-Graber-Jorgensen-2023.pdf>
- Dynan, K. E., Elmendorf, D., and Sichel, D. (2012). ‘The evolution of household income volatility’, *The B.E. Journal of Economic Analysis & Policy: Advances*, 12, 1–40. <https://doi.org/10.1515/1935-1682.3347>
- Farina, E., Green, C., and McVicar, D. (2020). ‘Zero-hours contracts and their growth’, *British Journal of Industrial Relations*, 58, 507–531. <https://doi.org/10.1111/bjir.12512>
- Farrell, D., and Greig, F. (2015). ‘Weathering volatility. Big data on the financial ups and downs of U.S. individuals’, JPMorgan Chase Institute. <https://www.jpmorganchase.com/content/dam/jpmc/jpmorgan-chase-and-co/institute/pdf/54918-jpmc-institute-report-2015-aw5.pdf>
- Farrell, D., Greig, F., and Yu, C. (2019). ‘Weathering Volatility 2.0’, JPMorgan Chase Institute. <https://www.jpmorganchase.com/content/dam/jpmc/jpmorgan-chase-and-co/institute/pdf/institute-volatility-cash-buffer-report.pdf>
- Ganong, P., Noel, P., Patterson, C., Vavra, J., and Weinberg, A. (2024). ‘Earnings instability’. Unpublished paper, University of Chicago. https://bpb-us-w2.wpmucdn.com/voices.uchicago.edu/dist/1/801/files/2024/03/earnings_instability-8ec8ba2640feedb4.pdf
- Gosselin, P. and Zimmerman, S. D. (2008). ‘Trends in income volatility and risk, 1970–2004’. Working Paper. Washington DC: Urban Institute. <https://www.urban.org/sites/default/files/publication/31731/411672-trends-in-income-volatility-and-risk-.pdf>
- Gottschalk, P. and Moffitt, R. (1994). ‘The growth of earnings instability in the U.S. labor market’, *Brookings Papers on Economic Activity*, 1, 217–254. <https://doi.org/10.2307/2534657>
- Gottschalk, P. and Moffitt, R. (2009). ‘The rising instability of U.S. earnings’, *Journal of Economic Perspectives*, 23, 3–24. <http://doi.org/10.1257/jep.23.4.3>
- Griffiths, R., and Wood, M. (2024). *Coping and Hoping. Navigating the ups and downs of monthly assessment in universal credit*. IPR Report, University of Bath. <https://www.bath.ac.uk/publications/coping-and-hoping-navigating-the-ups-and-downs-of-monthly-assessment-in-universal-credit/>

- Guvenen, F. (2009). ‘An empirical investigation of labor income processes’, *Review of Economic Dynamics*, 12, 58–79. <https://doi.org/10.1016/j.red.2008.06.004>
- Guvenen, F., Pistaferri, L., and Violante, G. L. (2022). ‘Global trends in income inequality and income dynamics: new insights from GRID’, *Quantitative Economics*, 13, 1321–1360. <https://doi.org/10.3982/QE2260>
- Hacker, J. S. (2018). ‘Economic security’. In: J. E. Stiglitz, J. P. Fitoussi, and M. Durand (eds), *For Good Measure: Advancing Research on Well-Being Metrics Beyond GDP*. Paris: OECD, 203–240. <https://doi.org/10.1787/9789264307278-en>
- Hacker, J. S. and Jacobs, E. (2008). ‘The rising instability of American family incomes, 1969–2004. evidence from the Panel Study of Income Dynamics’. Briefing Paper #213. Washington DC: Economic Policy Institute. <http://www.epi.org/content.cfm/bp213>
- Hacker, J. S., Huber, G. A., Nichols, A., Rehm, P., and Schlesinger, M. (2014). ‘The Economic Security Index: a new measure for research and policy analysis’, *Review of Income and Wealth*, 60, S5–S32. <https://doi.org/10.1111/roiw.12053>
- Hacker, J. S. and Rehm, P. (2022). ‘Reducing risk as well as inequality: assessing the welfare state’s insurance effects’, *British Journal of Political Science*, 52, 456–466. <https://doi.org/10.1017/S0007123420000034>
- Hannagan, A. and Morduch, J. (2016). ‘Income gains and month-to-month income volatility: evidence from the U.S. financial diaries’. In: *Economic Mobility: Research and Ideas on Strengthening Families, Communities and the Economy*. St. Louis, MO: Federal Reserve Bank of St. Louis. https://www.stlouisfed.org/community-development/publications/-/media/project/frbstl/stlouisfed/files/pdfs/community-development/econmobilitypapers/section2/econmobility_2-5hannaganmorduch_508.pdf
- Hardy, B. L. and Ziliak, J. P. (2014). ‘Decomposing trends in income volatility: the “wild ride” at the top and bottom’, *Economic Inquiry*, 52, 459–476. <https://doi.org/10.1111/ecin.12044>
- Hills, J., McKnight, A., and Smithies, R. (2006). Tracking income: how working families’ incomes vary through the year. Centre for Analysis of Social Exclusion Report 32. London: LSE. https://sticerd.lse.ac.uk/CASE/_NEW/PUBLICATIONS/abstract/?index=2310.

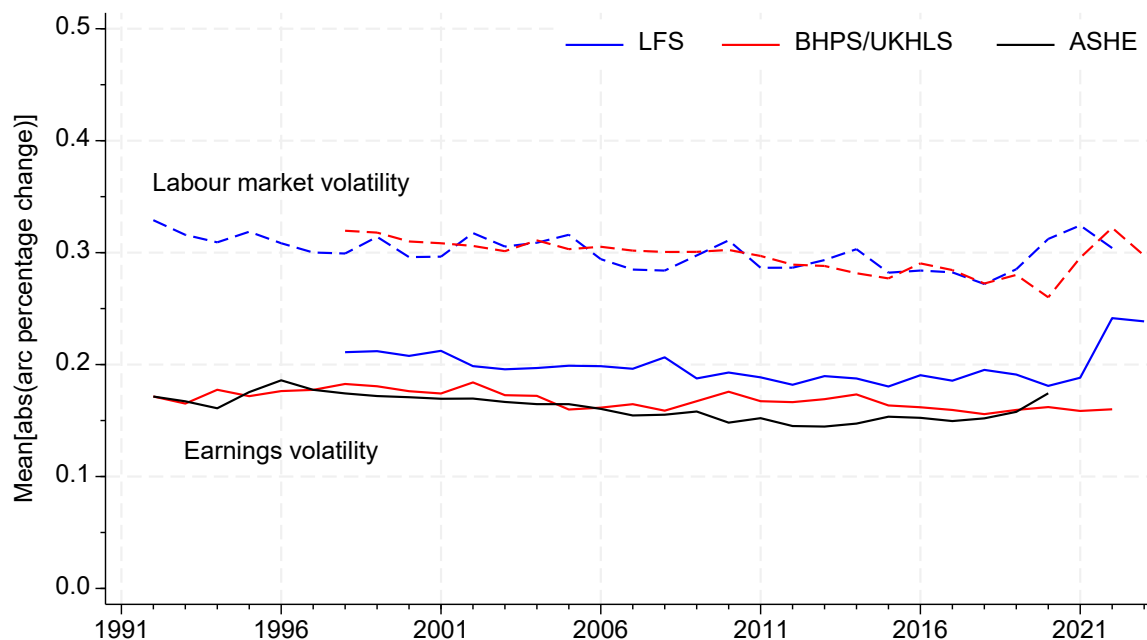
- Jenkins, S. P. (2011). *Changing Fortunes: Income Mobility and Poverty Dynamics in Britain*. Oxford: Oxford University Press.
<https://doi.org/10.1093/acprof:oso/9780199226436.001.0001>
- Jenkins, S. P. (2025). ‘Multiperiod income volatility measures’. Unpublished paper in preparation, LSE.
- Johnson, T.-M., Ghelani, D., Walker, R., and Charlesworth, Z. *Cheques and Imbalances. How income varies throughout the year for Universal Credit households*. London: Policy in Practice. https://policyinpractice.co.uk/wp-content/uploads/2025/03/Report_-_Cheques-and-imbalances-How-income-varies-throughout-the-year-for-Universal-Credit-households.-March-2025-1.pdf
- Jones, M. R., Bee, A., Eng, A., Houghton, H., Pharris-Ciurej, N., Porter, S. R., Rothbaum, J., and Voorheis, J. (2024). ‘Mobility, opportunity, and volatility statistics (MOVS): infrastructure files and public use data’. Working Paper CES 24-23, US Census Bureau. <https://www2.census.gov/library/working-papers/2024/adrm/ces/CES-WP-24-23.pdf>.
- Latner, J. (2018). ‘Income volatility and mobility: A conceptual exploration of two frameworks’. *Research in Social Stratification and Mobility*, 50, 50–63.
<https://doi.org/10.1016/j.rssm.2017.10.002>
- Menta, G., Wolff, E., and D’Ambrosio, C. (2023). ‘Income and wealth volatility: evidence from Italy and the U.S. in the past two decades’, *Journal of Economic Inequality*, 19, 293–313. <https://doi.org/10.1007/s10888-020-09473-4>
- Merfeld, J. D. and Morduch, J. (2024). ‘Poverty at higher frequency’.
<https://joshmerfeld.github.io/assets/papers/merfeldmorduch2024.pdf>.
- Millar, J., Cooke, K., and McLaughlin, E. (1989). ‘The employment lottery: risk and social security benefits’, *Policy and Politics*, 17, 75–81.
<https://doi.org/10.1332/030557389783219451>
- Moffitt, R., Abowd, A., Bollinger, C. R., Carr, M., Hokayem, C., McKinney, K., Weimers, E., Zhang, S., and Ziliak, J. (2023). ‘Reconciling trends in U.S. male earnings volatility: results from survey and administrative data’, *Journal of Business and Economic Statistics*, 41, 1–11. <https://doi.org/10.1080/07350015.2022.2102020>
- Moffitt, R. and Gottschalk, P. (1995). ‘Trends in the covariance structure of earnings in the U.S.: 1969–1987’. Unpublished paper.
https://www.econ2.jhu.edu/people/Moffitt/mg2_0795.pdf Later published in the

- Rediscovered Classics section, *Journal of Economic Inequality*, 2011, 9, 439–459.
<https://doi.org/10.1007/s10888-010-9154-z>
- Moffitt, R. and Gottschalk, P. (2012). ‘Trends in the transitory variance of male earnings: methods and evidence’, *Journal of Human Resources*, 47, 204–236.
<https://doi.org/10.3368/jhr.47.1.204>
- Morduch, J. and Schneider, R. (2017). *The Financial Diaries: How American Families Cope in a World of Uncertainty*. Princeton, NJ: Princeton University Press.
<https://doi.org/10.2307/j.ctvc77n3j>
- Morduch, J. and Siwicki, J. (2017). ‘In and out of poverty: episodic poverty and income volatility in the US financial diaries’, *Social Service Review*, 91, 390–421.
<https://doi.org/10.1086/694180>
- Morris, P. A., Hill, H. A., Gennetian, L. A., Rodrigues, C., and Wolf, S. (2015). ‘Income volatility in U.S. households with children: another growing disparity between the rich and the poor?’. Discussion Paper 1429-15, Institute for Research on Poverty.
<https://www.irp.wisc.edu/publications/dps/pdfs/dp142915.pdf>
- Murphy, L. (2021). CIPP payslip statistics comparison report 2008 – 2021. Chartered Institute of Payroll Professionals. <https://www.cipp.org.uk/static/40fc3f67-ed2c-41c3-a32d0460906697c4/Payslip-statistics-report-2021210705v1web.pdf>
- Nichols, A. and Rehm, P. (2014). ‘Income risk in 30 countries’, *Review of Income and Wealth*, 60, S98–S116. <https://doi.org/10.1111/roiw.12111>
- OECD (2023). *On Shaky Ground? Income Instability and Economic Insecurity in Europe*, Paris: OECD Publishing. <https://doi.org/10.1787/9bffe6a6-en>
- Office for National Statistics (2017). *Average Weekly Earnings, bonus payments in Great Britain: 2017*, Statistical Bulletin.
<https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinhours/bulletins/averageweeklyearningsbonuspaymentsingreatbritain/2017>.
- Osberg, L. (2015). ‘How should one measure economic insecurity?’ Statistics Working Paper 2015/01. Paris: OECD. <https://dx.doi.org/10.1787/5js4t78q9lq7-en>
- Ostrovsky, Y. (2012). ‘The correlation of spouses’ permanent and transitory earnings and family earnings inequality in Canada’, *Labour Economics*, 19, 756–768.
<https://doi.org/10.1016/j.labeco.2012.07.005>
- Pew Trusts (2015). *Americans’ Financial Security. Perception and Reality*.
https://www.pewtrusts.org/-/media/assets/2015/02/fsm-poll-results-issue-brief_artfinal_v3.pdf.

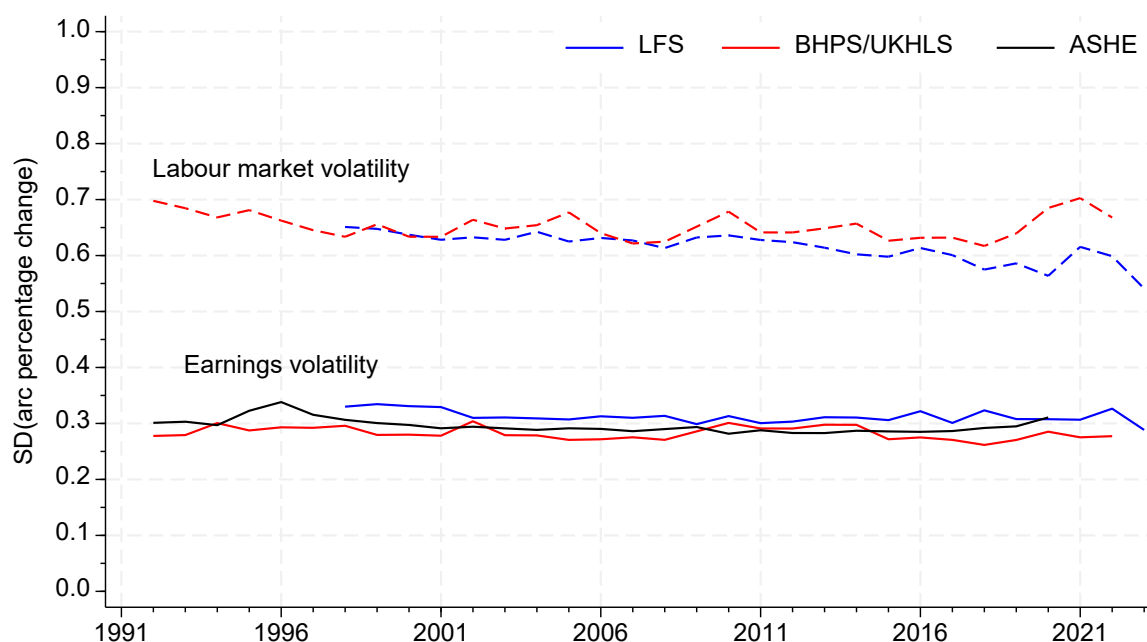
- Romich, J. and Hill, H. D. (2017), ‘Income instability and income support programs: recommendations for policy and practice’, Family Self-Sufficiency and Stability Research Consortium, Mathematica. <https://mathematica.org/publications/income-instability-and-income-support-programs-recommendations-for-policy-and-practice>.
- Sabelhaus, J. and Song, J. (2010). ‘The Great Moderation in micro labor earnings’, *Journal of Monetary Economics*, 57, 391–403. <https://doi.org/10.1016/j.jmoneco.2010.04.003>
- Shin, D., and Solon, G. (2011). Trends in men’s earnings volatility: what does the Panel Study of Income Dynamics show?, *Journal of Public Economics*, 95, 973–982. <https://doi.org/10.1016/j.jpubeco.2011.02.007>
- Shiu, J.-L., Zhang, S., and Gottschalk, P. (2024). ‘Family income dynamics 1970-2018: putting the pieces together’, *Journal of Labor Economics*, forthcoming. <https://doi.org/10.1086/732769>.
- Shorrocks, A. F. (1978). ‘Income inequality and income mobility’, *Journal of Economic Theory*, 19, 376–393. [https://doi.org/10.1016/0022-0531\(78\)90101-1](https://doi.org/10.1016/0022-0531(78)90101-1)
- Tomlinson, D. (2018). *Irregular Payments: Assessing the breadth and depth of month to month earnings volatility*. London: Resolution Foundation. <https://www.resolutionfoundation.org/publications/irregular-payments/>.
- WED (2024). *HMRC PAYE & SA data. Quick guide for users*, August. <https://www.wagedynamics.com/wp-content/uploads/2024/07/HMRC-Quick-User-Guide-.pdf>.
- Western, B., Bloome, D., Sosnaud, B., and Tach, L. M. (2016). ‘Trends in income insecurity among U.S. children, 1984–2010’, *Demography*, 53, 419–447. <https://doi.org/10.1007/s13524-016-0463-0>
- Winship, S. (2011). ‘Economic instability trends and levels across household surveys’. Working Paper 11-13, National Poverty Centre, University of Michigan. https://npc.umich.edu/publications/working_papers/publication_id/209/.
- Ziliak, J. P., Hardy, B., and Bollinger, C. R. (2011). ‘Earnings volatility in America: evidence from matched CPS’, *Labour Economics*, 18, 742–754. <https://doi.org/10.1016/j.labeco.2011.06.015>
- Ziliak, J. P., Hokayem, C., and Bollinger, C. R. (2023). ‘Trends in earnings volatility using linked administrative and survey data’. *Journal of Business and Economic Statistics*, 41, 12–19. <https://doi.org/10.1080/07350015.2022.2102023> 99999

Figure 1. Volatility trends (monthly earnings a year apart): UK survey-based estimates

(a) Mean absolute arc percentage change

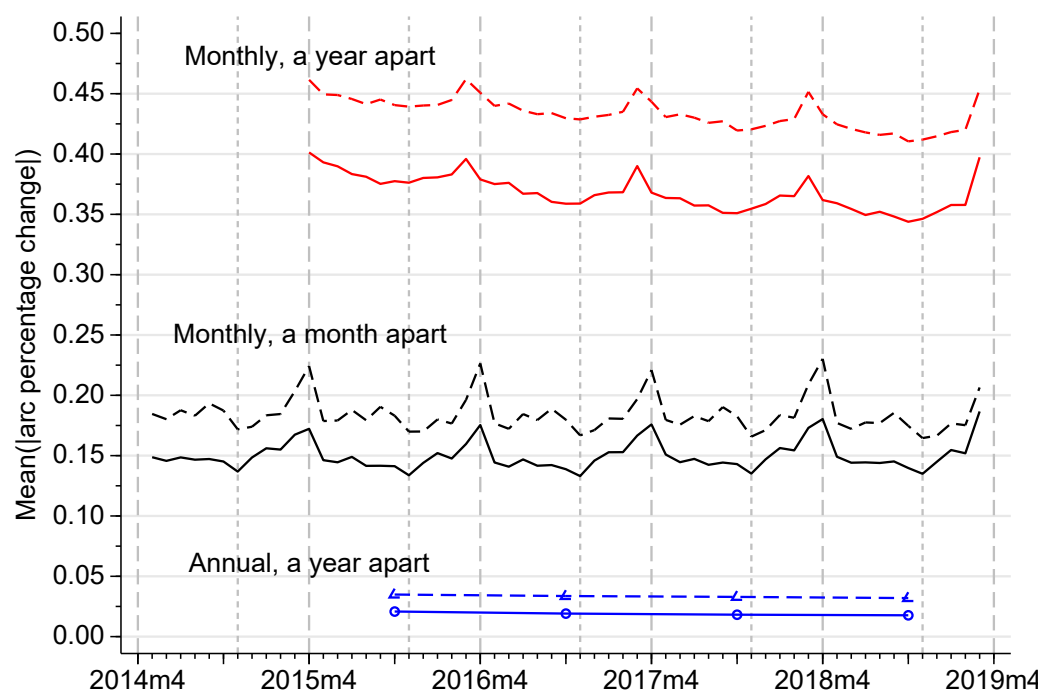


(b) Standard deviation of arc percentage changes



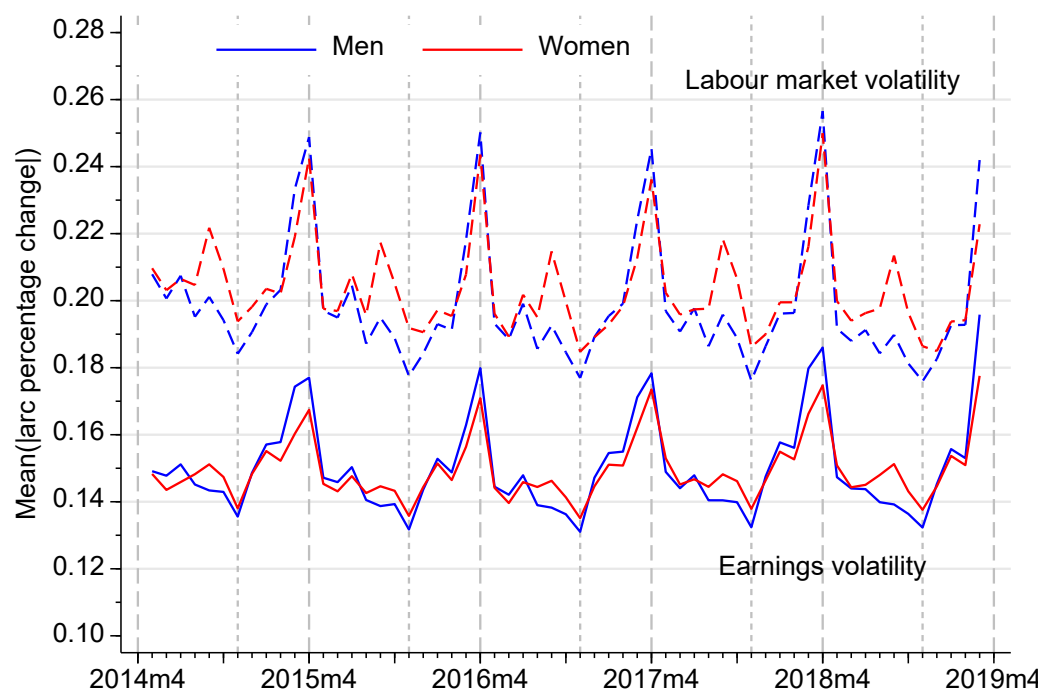
Authors' calculations from the quarterly Labour Force Survey (LFS), British Household Panel Survey (BHPS), UK Household Longitudinal Survey (UKHLS), and Annual Survey of Hours and Employment (ASHE). ASHE estimates refer to the April of each year. Otherwise, estimates shown are for yearly averages rather than for each interview month or quarter, because of small sample sizes at sub-annual levels for the BHPS, UKHLS, and LFS.

Figure 2. Pay volatility: changes in monthly earnings a month apart, a year apart, and change in annual earnings



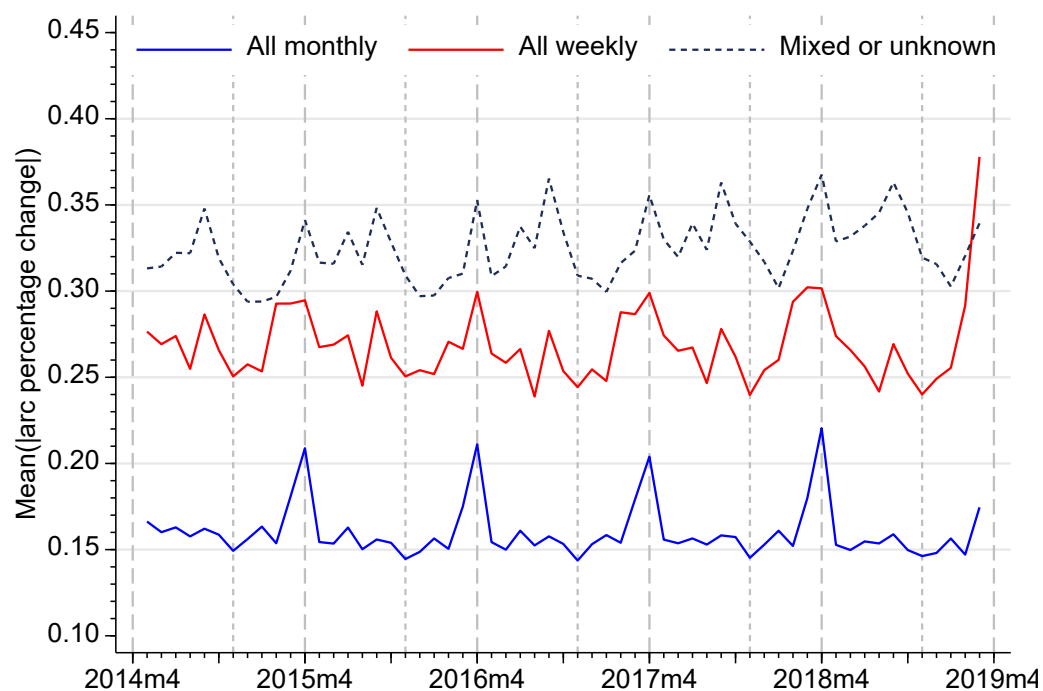
Source: HMRC PAYE data, HMRC file. Notes. $Mn(|a_t|)$ is the mean absolute arc percentage change in earnings. Dashed lines in series refer to labour market volatility (i.e., including workers with entries to or exits from positive earnings). Solid lines refer to earnings volatility (among workers with positive earnings at $t-1$ and t). Vertical dashed lines identify April and November. Horizontal axis labels refer to time t for changes between $t-1$ and t . Annual estimates refer to earnings over a financial year.

Figure 3. Monthly earnings volatility, by sex



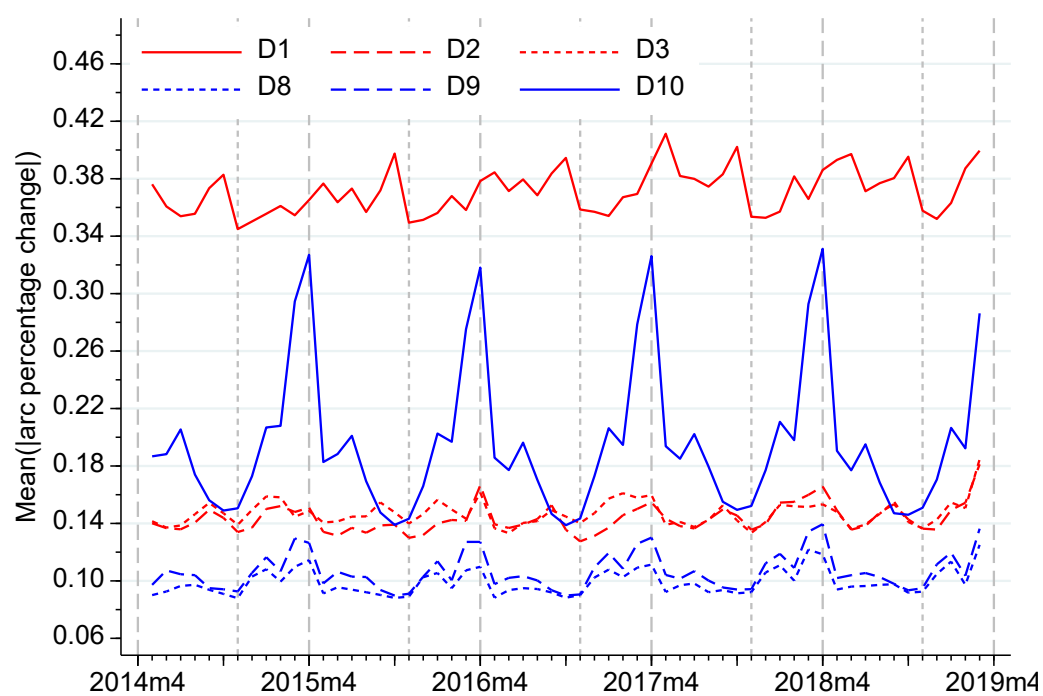
Source: HMRC PAYE data, HMRC file. Notes: as for Figure 2.

Figure 4. Monthly earnings volatility, by payslip type



Source: HMRC PAYE data, HMRC file. Notes: as for Figure 2. Fraction of person-month observations that refer to payslips of monthly type is 73%, of weekly type 23%, of mixed or unknown type 3%. Estimates refer to observations with positive earnings in months $t-1$ and t .

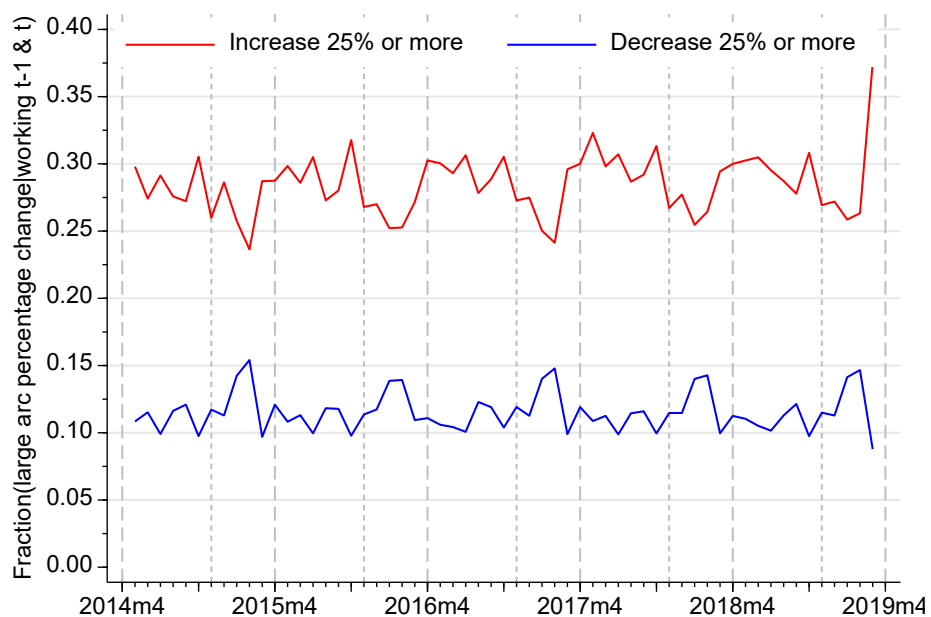
Figure 5. Monthly earnings volatility, by decile group of earnings



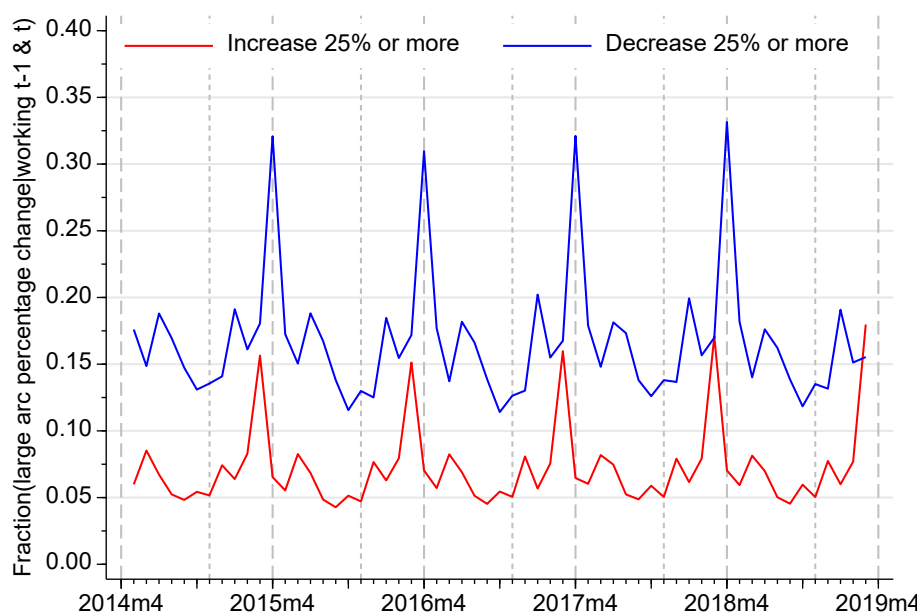
Source: HMRC PAYE data, HMRC file. Notes: as for Figure 2. Estimates refer to observations with positive earnings in months $t-1$ and t . Decile groups refer to the distribution of earnings in month $t-1$. D1 is the bottom tenth, D10 is the top tenth.

Figure 6. Fractions with monthly earnings increase and decrease of at least 25%, by decile group of earnings

(a) Bottom decile group of earnings

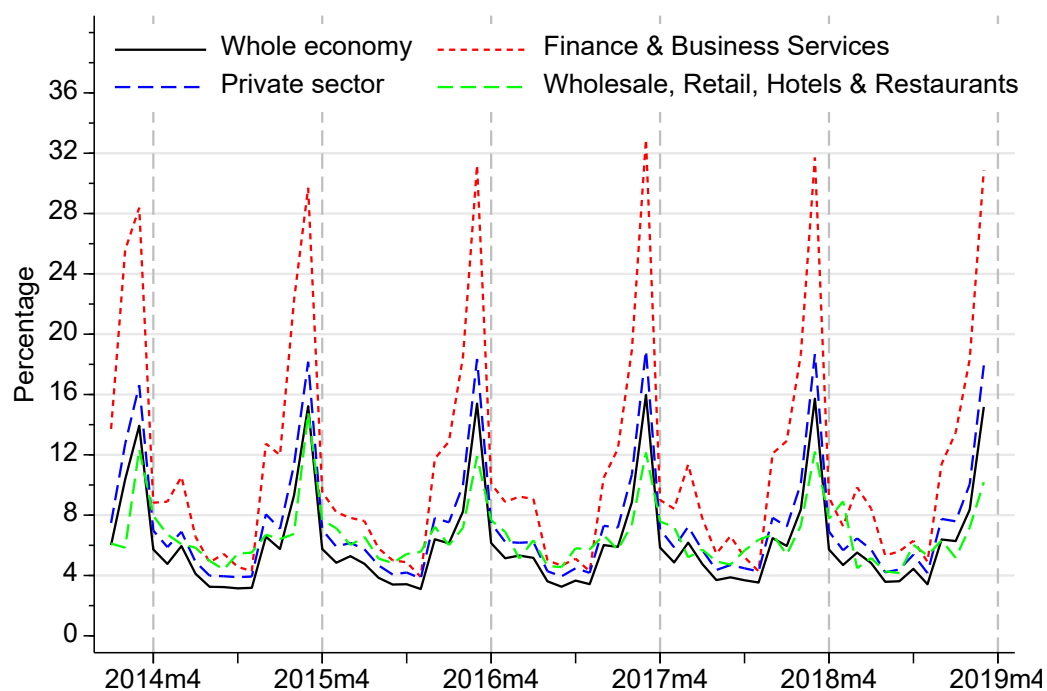


(b) Top decile group of earnings



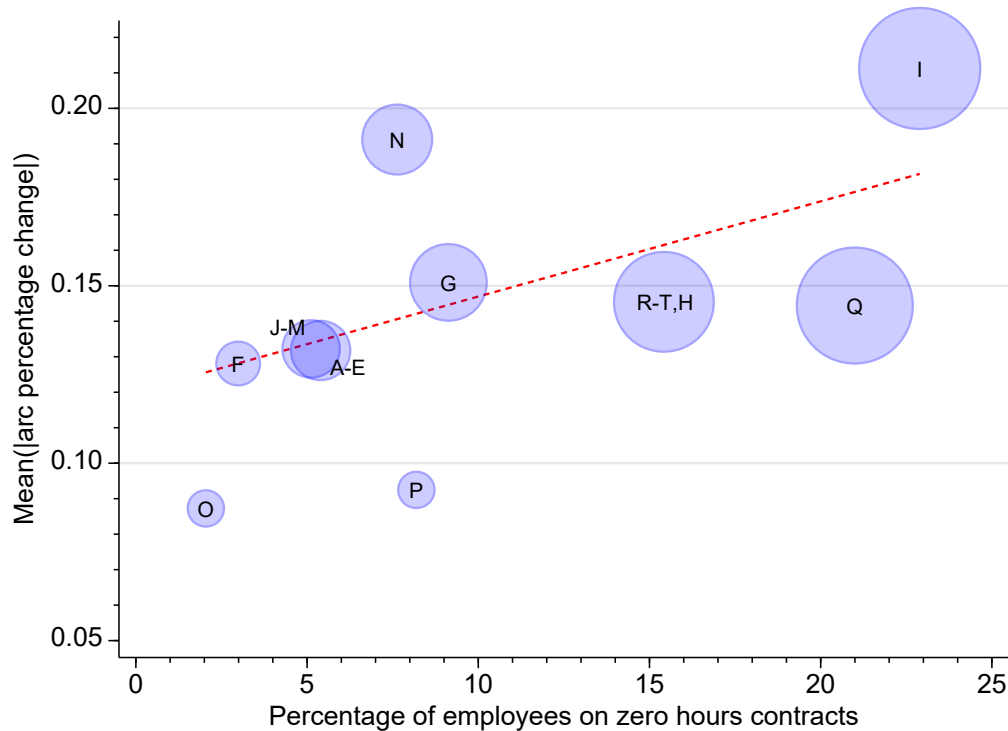
Source: HMRC PAYE data, HMRC file. Notes. Estimates refer to observations with positive earnings in months $t-1$ and t . Decile groups of earnings refer to the monthly earnings distribution for month $t-1$.

Figure 7. Bonuses as percentage of average weekly earnings, UK, by month and sector



Source: Source: ONS (2025), Dataset EARN02 Average Weekly Earnings by Sector - Not Seasonally Adjusted (spreadsheet at <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/averageweeklyearningsbysectorearn02>).

Figure 8. Earnings volatility is higher in industries in which employers make greater use of zero-hours contracts



Source: authors' calculations from (i) ONS (17 February 2025), EMP17: People in employment on zero-hours contracts (spreadsheet at <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/datasets/emp17peopleinemploymentonzerohourscontracts/current>), and (ii) HMRC PAYE data, Linked-ASHE file.

Notes. Bubbles show estimates for SIC industry sections: A–E Production including agriculture; F Construction; G Wholesale and retail; I Accommodation and food; J–M Information, finance, professional; N Administrative and support services; O Public administration; P Education; Q Health and social work; R–T, H Transport, arts, other services. Zero-hours contract proportions are derived from the Labour Force Survey and refer to employees aged 16+ years. Displayed are averages of quarterly estimates for 2014–2018. Bubble size is proportional to the total number of employees in industry. Volatility estimates refer to those working in months $t-1$ and t . The dashed red line is a least-squares regression line fit to the bubble plot.

Table 1. Month-to-month earnings volatility, by subgroup and volatility measure (all working at $t-1$ and t)

	Size (% of total)	Arc percentage change		Arc percentage change			
		Mean	%	SD	%	Variance	%
All (all working at $t-1$ and t)	100.00	0.136	100	0.286	100	0.082	100
<i>Sex</i>							
Female	48.52	0.137	101	0.290	101	0.084	102
Male	51.48	0.134	99	0.282	99	0.080	97
<i>Age (years)</i>							
20–24	12.53	0.186	137	0.348	122	0.121	148
25–29	16.58	0.146	107	0.296	103	0.088	107
30–34	16.18	0.137	101	0.288	101	0.083	101
35–39	13.28	0.129	95	0.279	98	0.078	95
40–44	11.39	0.123	90	0.270	94	0.073	89
45–49	11.52	0.117	86	0.259	91	0.067	82
50–54	10.45	0.114	83	0.255	89	0.065	79
55–59	8.08	0.116	85	0.262	91	0.068	83
<i>Decile group of pay (at $t-1$)</i>							
D1 (lowest)	10.00	0.283	208	0.494	173	0.244	297
D2	10.00	0.155	114	0.290	101	0.084	102
D3	10.00	0.128	94	0.248	87	0.061	75
D4	10.00	0.107	79	0.219	77	0.048	58
D5	10.00	0.100	74	0.207	73	0.043	52
D6	10.00	0.096	71	0.203	71	0.041	50
D7	10.00	0.095	70	0.202	71	0.041	50
D8	10.00	0.095	70	0.205	72	0.042	51
D9	10.00	0.104	76	0.225	79	0.051	62
D10 (highest)	10.00	0.192	141	0.381	133	0.145	177
<i>Low-paid occupation</i>							
Non-low paid occupations	68.57	0.121	89	0.274	96	0.075	91
Agriculture	0.62	0.135	99	0.274	96	0.075	92
Food processing	1.40	0.163	120	0.295	103	0.087	106
Textiles	0.22	0.139	102	0.257	90	0.066	81
Retail	8.27	0.162	119	0.303	106	0.092	112
Hospitality	4.70	0.197	145	0.357	125	0.127	155
Cleaning	2.66	0.166	122	0.309	108	0.096	117
Social care	2.88	0.205	151	0.349	122	0.122	148
Childcare	1.46	0.155	114	0.301	105	0.091	110
Leisure	0.85	0.187	137	0.347	121	0.120	147
Hairdressing	0.67	0.129	95	0.264	92	0.069	85
Office work	1.66	0.131	97	0.274	96	0.075	92
Non-food processing	1.63	0.159	117	0.291	102	0.084	103
Storage	2.58	0.162	119	0.300	105	0.090	110
Transport	1.82	0.130	96	0.262	92	0.069	84
<i>High-paid occupation</i>							
Not high-paid	86.49	0.136	100	0.282	99	0.080	97
Functional managers & directors	2.62	0.136	100	0.346	121	0.120	146
Fin. institutions managers & directors	0.25	0.177	130	0.378	132	0.143	174
IT, telecoms	2.45	0.108	79	0.265	93	0.070	86
Business, research & admin professionals	2.29	0.110	81	0.282	99	0.080	97
Business, finance, etc., assoc professionals	1.93	0.127	93	0.302	106	0.091	111

Sales & marketing & related	3.97	0.155	114	0.326	114	0.106	130
<i>Pay slip type</i>							
All monthly	74.82	0.118	87	0.265	93	0.070	86
All weekly	22.27	0.179	132	0.330	116	0.109	133
Mixed or unknown	2.91	0.248	182	0.411	144	0.169	206
<i>Whether has 1+ jobs</i>							
One job	92.24	0.127	94	0.274	96	0.075	92
More than one job	7.76	0.233	171	0.401	140	0.161	196
<i>Contract type</i>							
Permanent	93.20	0.129	95	0.274	96	0.075	92
Temporary	6.80	0.227	167	0.412	144	0.170	207
<i>Hourly paid or not</i>							
Not hourly paid	60.38	0.116	85	0.269	94	0.073	88
Hourly paid	39.62	0.166	122	0.310	108	0.096	117
<i>Whether working full- or part-time</i>							
Part-time	25.65	0.171	126	0.332	116	0.110	134
Full-time	74.35	0.123	91	0.268	94	0.072	88
<i>Public/private sector job</i>							
Private	82.91	0.144	106	0.295	103	0.087	106
Public	17.09	0.097	71	0.238	83	0.056	69
<i>Same job more than one year</i>							
Same job 1+ years	79.32	0.126	93	0.270	94	0.073	89
Not same job	20.68	0.173	127	0.340	119	0.115	141

Source: HMRC PAYE data, Linked-ASHE file. Notes: Person-month observations pooled over 60 months. Sample is observations with positive earnings in months $t-1$ and t . All estimates are weighted. ‘%’: subgroup estimate as % of corresponding All estimate (row 1). Subgroup characteristics, except pay frequency type and multiple jobholding, are derived from ASHE, and refer to April of each financial year. Pay decile groups refer to decile groups of the Linked-ASHE file sample. High-paid subgroups correspond to SOC2010 groups 113x, 115x, 213x, 242x, 353x, and 354x, respectively.

Online Supplementary Material

What do we know about income and earnings volatility?

Mike Brewer, Nye Cominetti, and Stephen P. Jenkins

24 March 2025

Appendices to accompany Section 5 of the main text

Appendix A. Construction of consistent real monthly pay variables

Appendix B. Volatility time series charts for subgroups

Appendix C. Multiperiod volatility, with subgroup breakdowns

Appendix D. Volatility regressions and volatility profiles for selected person types

Appendix E. Comparisons of volatility distributions and volatility breakdowns: HMRC versus Linked-ASHE file estimates (time series charts and tables)

Appendix A. Construction of consistent real monthly pay variables

Here we discuss how we created a consistent monthly pay variable and compare our approach with some potential alternatives. We also discuss how we adjusted observed pay variables for inflation over our 60-month sample period.

Consistent monthly pay variables for workers with different pay periods

In the UK's PAYE system, employers can choose their pay frequency but must inform the tax authorities (HMRC) what the pay period is. There are many possibilities: weekly or some other weekly type (e.g., fortnightly, four-weekly), monthly or some other monthly type (e.g., two-monthly), or at some other frequency (e.g., annual).

To put the data associated with different pay periods on a common monthly basis, first the WED (2024) team created two longitudinal files: (a) a monthly file containing all payslips that are monthly or some other monthly type (around 14 million payslips), and (b) a weekly file in which all payslips are weekly or some weekly type (around 15 million payslips). Payslips referring to the payment of occupational (employer administered) pensions were excluded. Second, where the payslips for a worker referred to multiples of months (in file a) or multiples of weeks (in file b), the team spread pay evenly over the relevant period, yielding monthly and weekly pay amounts pro rata respectively. Third, for the weekly cases in file (b), because months have different numbers of weeks (and hence some weeks span different calendar months), the weekly amounts derived at step 2 were converted to average daily amounts ('daily-ised') and these were summed to monthly totals over the days within each of the 60 calendar months. Fourth, also for the weekly payslip types in file (b), because years have 365 or 366 days, there are one or two extra days 'left over' from step three (52×7 days per week = 364 days). For the weekly paid, the left-over pay was allocated to week 52 of each financial year; for the fortnightly paid, it was spread over weeks 51 and 52; for the four-weekly paid, it was spread over weeks 49–52.

This method of producing consistent monthly pay estimates means that, for weekly cases in file (b), derived monthly pay varies according to the length of the calendar month even if a weekly-paid employee receives the same pay amount each week. For example, a non-leap year February has around 10% fewer days than January, and March has around 11% more days than February. However, the effects of this are muted in our analysis by two factors. One is our use of the arc percentage change (because its numerator averages pay at $t-1$ and t) in summary measures such as $Mn(a_t)$ and $SD(a_t)$. The second is that when we

document the prevalence of ‘large’ income changes, we focus on the prevalence of changes of at least 25%.

This type of variation in our derived monthly pay variable is a feature not a bug. Someone who is paid weekly really does get paid more in March than in February. This is a non-trivial point because many household bills (rent, utilities) are charged by the calendar month and the UK’s principal cash transfer for non-working and working families, Universal Credit is assessed monthly on monthly income and paid monthly in arrears. (More about interactions between pay period and Universal Credit reference periods in Section 6 of the main text.)

There are other approaches to deriving a consistent monthly pay variable, as follows. First, one could pretend all months have the same length. That is, we could adjust weekly-paid data as described above but then multiply the derived monthly earnings by the number of days in the month times the number of days in the financial year divided by 12. In this approach, individuals with unchanging weekly earnings would then have unchanging ‘monthly’ earnings. However, the derived variable for earnings in March for someone paid weekly would be less than what they actually got paid.

Second, one could define total earnings per month to equal total earnings received in the month divided by the number of pay packets per month. Ganong et al. (2024) apply this rule to their US data. The approach implies that individuals with unchanging weekly earnings would have unchanging ‘monthly’ earnings. However, it also ignores the fact that total earnings in a month for a weekly-paid worker can vary in reality.

A third alternative would be to divide the year into 13×4 -week periods, which is the approach taken by Hills et al. (2006) whose data covered 52 weeks in total. The problem with this approach is that it also requires assumptions about how to allocate 12 months of pay to the 13 periods in each year. See Hills et al. (2006, p. 27) for a discussion of their approach in this situation.

Overall, all approaches to creating a consistent monthly earnings variable for workers with different length pay periods have some undesirable features. Our view is that the approach we have taken is the best in the context of our study.

Derivation of real monthly earnings

We converted nominal monthly earnings to real purchasing power terms in March 2019 prices using the monthly Consumer Prices Index including owner occupiers’ housing costs

(‘CPIH’). All previous studies of earnings volatility taking an annual perspective have adjusted earnings to account for inflation, including Bell et al.’s (2022) contribution to the GRID project. Among the few studies of volatility based on monthly income and earnings, both real and nominal amounts are used. Bania and Leete (2009), for example, use real incomes, but adjustments for inflation are not mentioned by Hills et al. (2006), Hannigan and Morduch (2015), Farrell and Grieg (2015), Farrell et al. (2019), Ganong et al. (2024), or Gruedahl et al. (2023). Tomlinson (2018) explicitly states that he adjusts earnings using the CPIH when examining annual changes in pay but otherwise works with nominal earnings. His reasoning is that “monthly changes in inflation will just be too small to have a bearing on the lived experience of pay change over such a short time period. In addition, it is likely that families think of changes in budgets from month to month in cash rather than price-adjusted terms.” (2018, 25). We agree with Tomlinson’s first point. Whether households think in nominal or real terms is more debatable and the point loses force in any case because we are following workers over five years, not one.

In our article we present volatility estimates based on real monthly earnings. We have also repeated our main analyses using nominal monthly earnings. The main impact is on the estimated fraction of workers with zero (arc percentage) change in earnings between month $t-1$ and t . Pooling the data for men and women and all 60 months, and focusing on those working at $t-1$ and t , the fraction with exactly zero change in real monthly earnings is 4.2% but 31.4% for nominal monthly earnings (see Table A1 below). However, otherwise, the distributions of real and nominal earnings changes, and subgroup differences, are very similar.

Figures A1 and A2, comparing relative frequency distributions and empirical cumulative distribution functions, demonstrate the similarity of the distributions of arc percentage changes for real and nominal earnings. From Table A1, we see that, looking at all individuals (i.e., including those moving from non-employment to employment in consecutive months, or vice versa), the median change for real earnings is 0.000 compared to 0.000 for nominal earnings. The respective means are 0.003 and 0.004, standard deviations 0.461 and 0.461, skewness -0.029 and 0.024 , and kurtosis 14.481 and 14.485 . Looking at individuals working in months $t-1$ and t , corresponding statistics are also remarkably similar.

In so far as comparisons are possible, we would note that the fraction of (pooled) person-month observations with exactly zero change in nominal earnings in our data is the same as reported by Ganong et al. (2024, Table 1), 0.31. The $SD(a_i)$ for those working in two consecutive months are somewhat different, 0.461 for the UK, and 0.28 for the USA. For

Denmark, the fraction with exactly zero change in monthly earnings is around 45% (estimated from Druedahl et al., 2023, Figure 1.1).


The reason differences between real and nominal earnings change distributions are negligible is because monthly CPIH inflation rates over our sample period were small, never more than $\pm 0.5\%$, and less if January is excluded. For the same reasons, using real and monthly earnings provide very similar pictures of the seasonal pattern of volatility within each year, and of volatility differences across subgroups of workers. Figure A3 shows the CPI monthly inflation rate over the sample period and adjacent years.

<Figures A1, A2, and A3 near here>


<Table A1 near here>

Figure A1. Relative frequencies (%) by band of arc percentage change in monthly earnings

(a) Real earnings (HMRC file), as in main text




(b) Nominal earnings (HMRC file)




Source: HMRC PAYE data, HMRC file, 60 months pooled. Note: bands are not of equal width. '(x,y]': greater than x and less than or equal to y.

Figure A2. Empirical cumulative distribution functions (ECDFs) for arc percentage changes in monthly earnings

(a) Real earnings

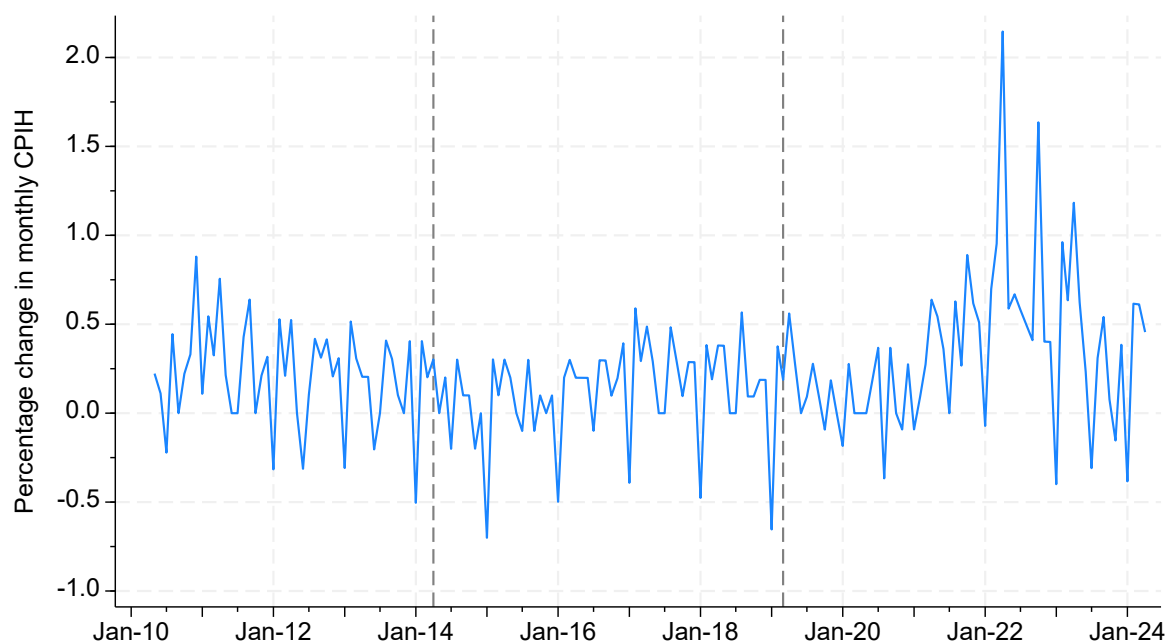


(b) Nominal earnings



Source: HMRC PAYE data, HMRC file, 60 months pooled. Notes: each ECDF shows the estimated fractions of individuals with an arc percentage change less than or equal to a specific arc percentage level (shown on horizontal axis). ECDFs estimated using 100 equally sized bins over the range between zero and one.

Figure A3. CPIH monthly inflation rates, April 2010 to April 2024



Source: UK Office for National Statistics, [Consumer Prices Index including owner occupiers' housing costs](#). Note: the vertical dashed lines demarcate the sample period (April 2014 to March 2019).

Table A1. Summary statistics for distributions on arc percentage changes in earnings

(a) Real earnings (as used in main text)

Sample	%	Mean	SD	p_{50}	p_{10}	p_{25}	p_{75}	p_{90}	Skewness	Kurtosis	Fraction with zero
All	100.00	0.003	0.461	0.000	-0.192	-0.012	0.022	0.205	-0.029	14.481	0.222
All working at $t-1$ and t	76.48	0.002	0.316	-0.002	-0.209	-0.036	0.042	0.219	-0.031	14.784	0.042

(b) Nominal earnings

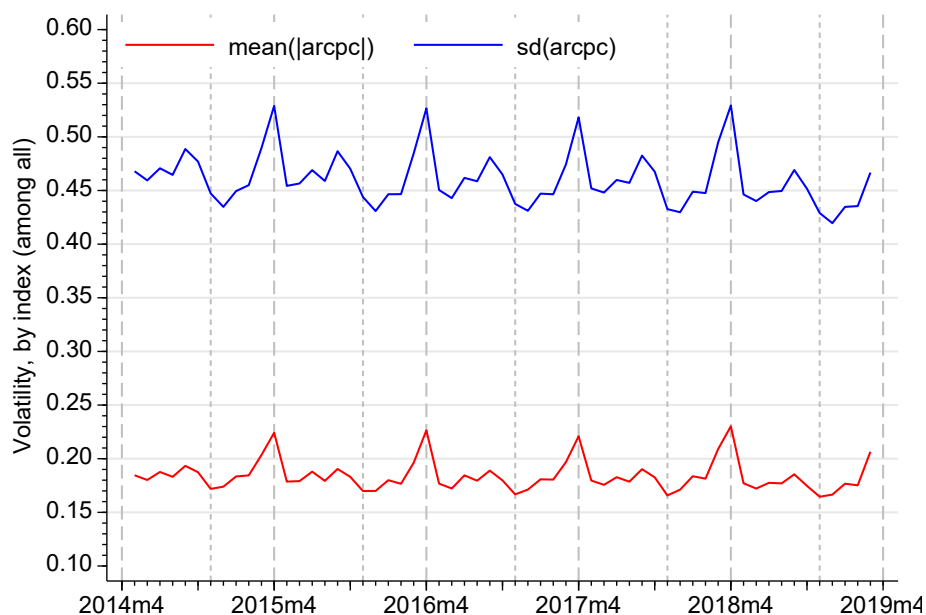
Sample	%	Mean	SD	p_{50}	p_{10}	p_{25}	p_{75}	p_{90}	Skewness	Kurtosis	Fraction with zero
All	100.00	0.004	0.461	0.000	-0.191	-0.011	0.023	0.206	0.024	14.485	0.432
All working at $t-1$ and t	76.48	0.004	0.316	0.000	-0.207	-0.034	0.044	0.221	-0.025	14.807	0.314

Source: HMRC PAYE data, HMRC file, 60 months pooled. Note: number of person-months = 15,859,140 (all); 12,129,585 (all working in months $t-1$ and t).

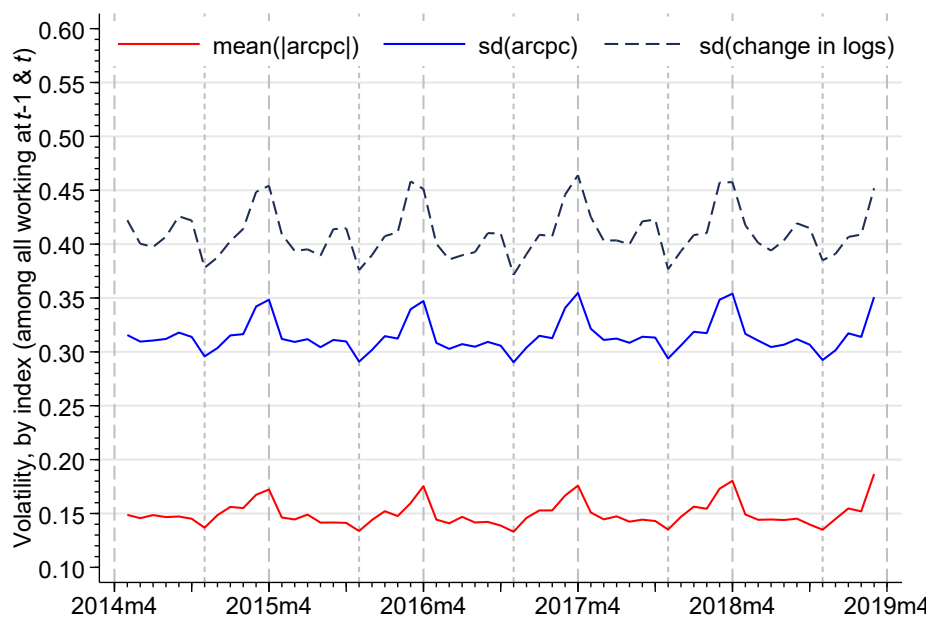
Appendix B. Additional volatility time series charts for subgroups

Figure B1. Different volatility indices show the same patterns of seasonality and overall trends

(a) Changes in monthly earnings, one month apart: all workers

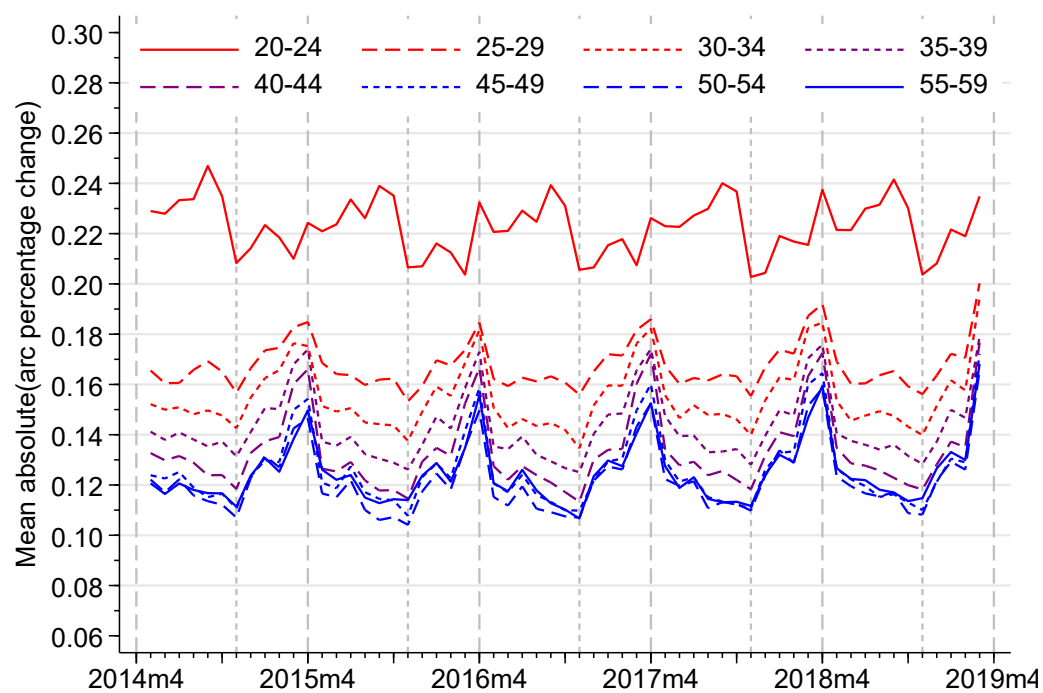


(b) Changes in monthly earnings, one month apart: all with positive earnings at $t-1$ and t



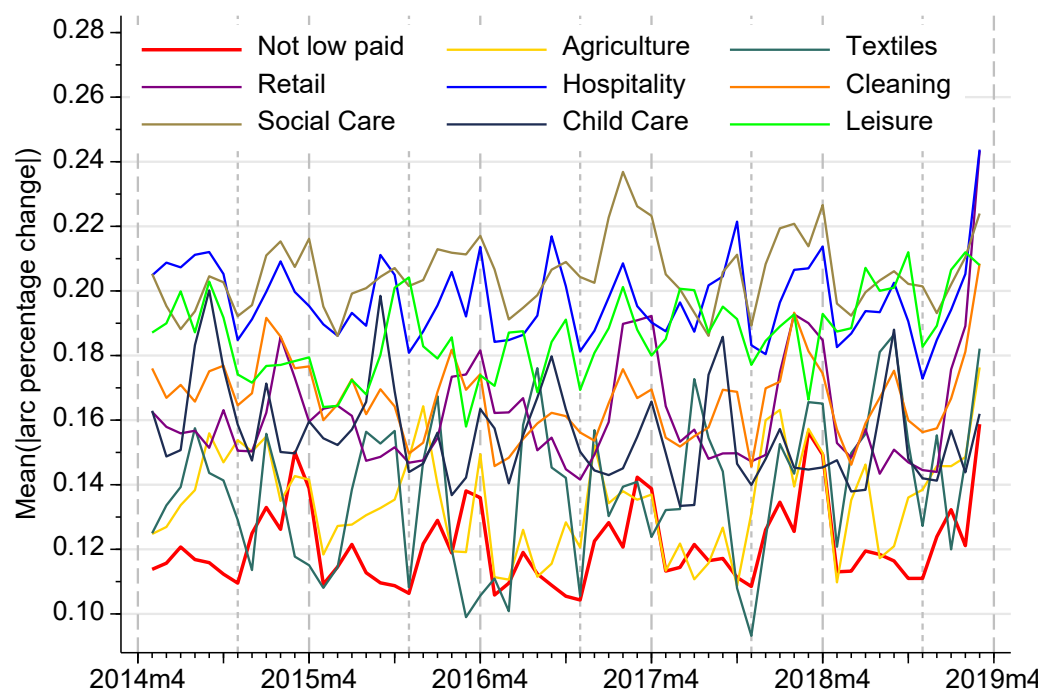
Source: HMRC PAYE data, HMRC file. Notes. Log earnings changes are undefined if earnings equal zero, so the SD(change in logs) measure is not shown in panel (a).

Figure B2. Monthly pay volatility, by age group



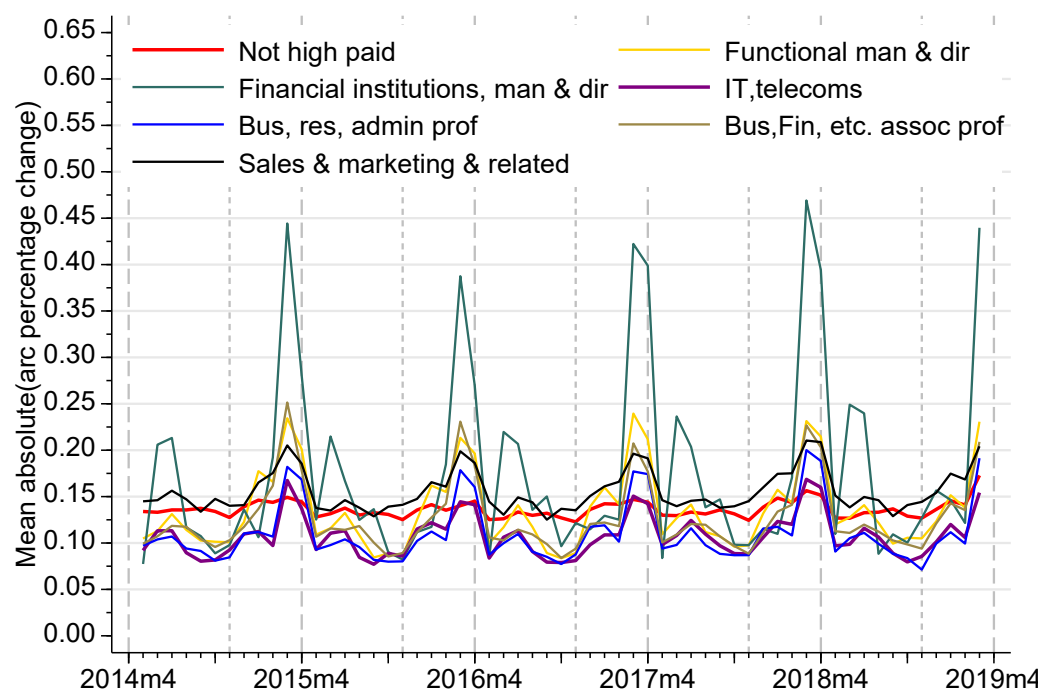
Source: HMRC PAYE data, HMRC file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t . Legend refers to age group (in years).

Figure B3. Monthly pay volatility, by selected low-paid occupation



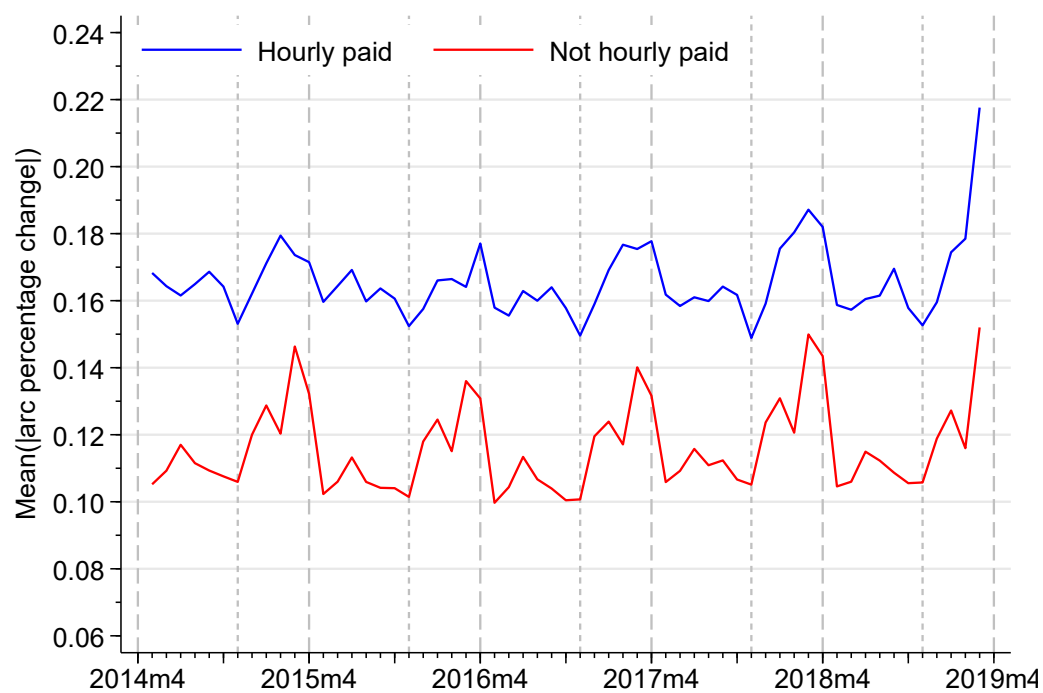
Source: HMRC PAYE data, Linked-ASHE file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t .

Figure B4. Monthly pay volatility, by selected high-paid occupation



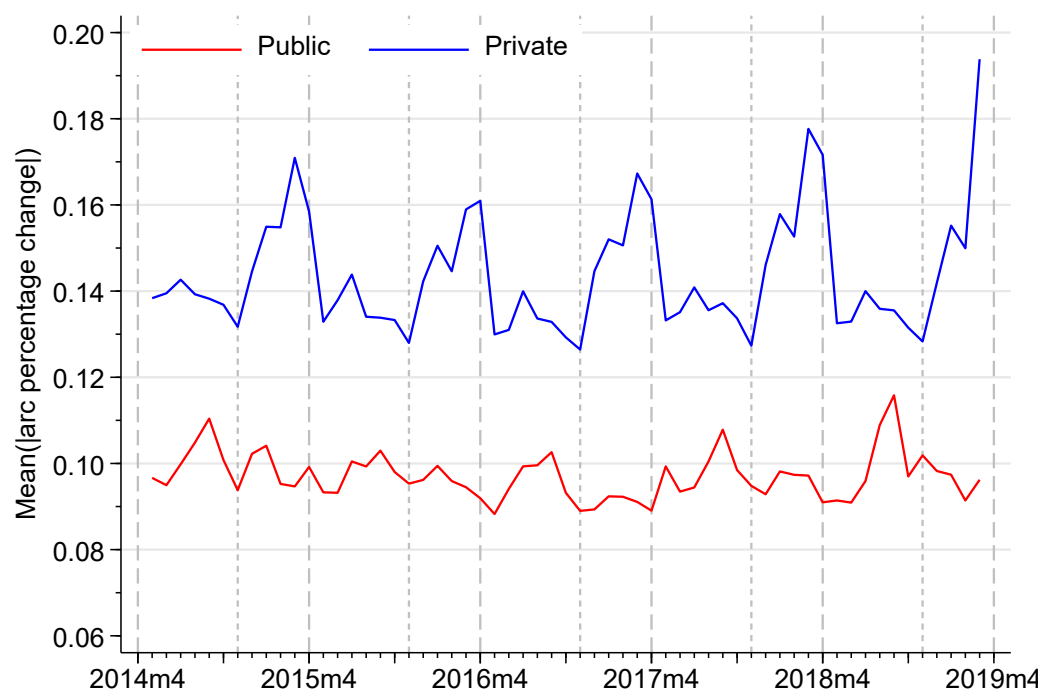
Source: HMRC PAYE data, Linked-ASHE file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t . High-paid subgroups correspond to SOC2010 groups 113x, 115x, 213x, 242x, 353x, and 354x, respectively (see main text).

Figure B5. Monthly pay volatility, by whether hourly paid or not



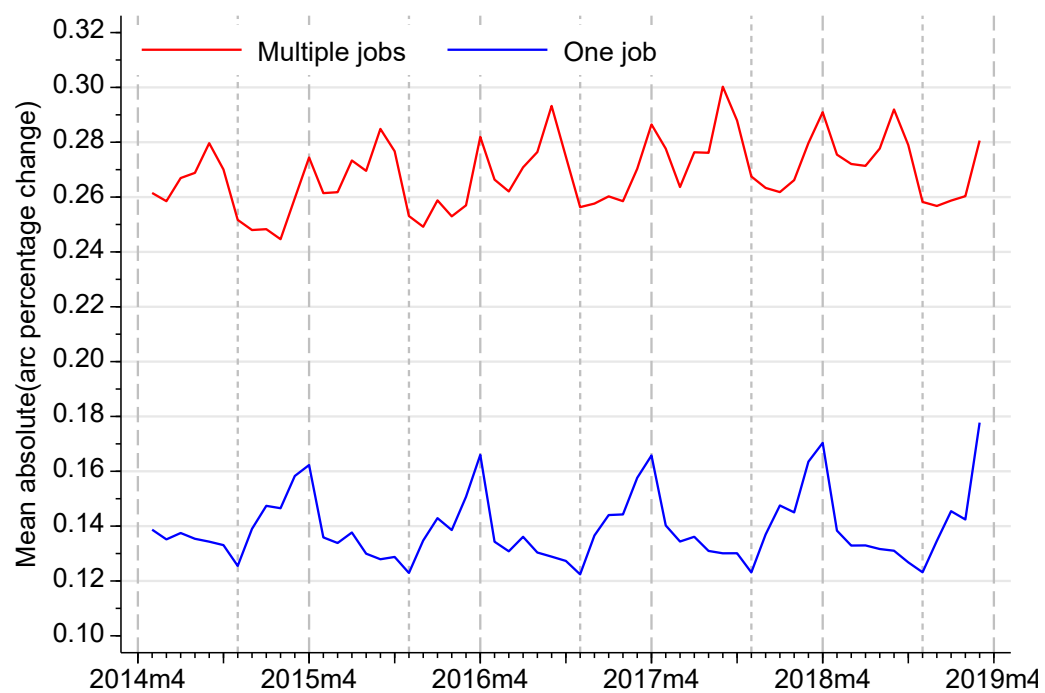
Source: HMRC PAYE data, Linked-ASHE file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t .

Figure B6. Monthly pay volatility, by whether works in the private or public sector



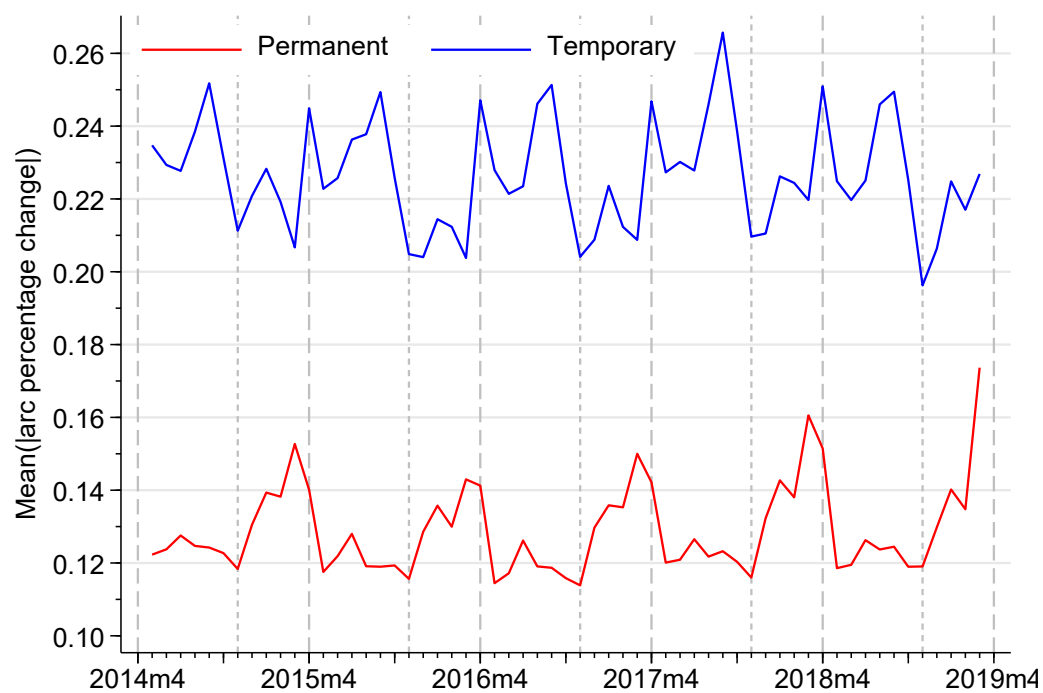
Source: HMRC PAYE data, Linked-ASHE file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t .

Figure B7. Monthly pay volatility, by whether has more than one job



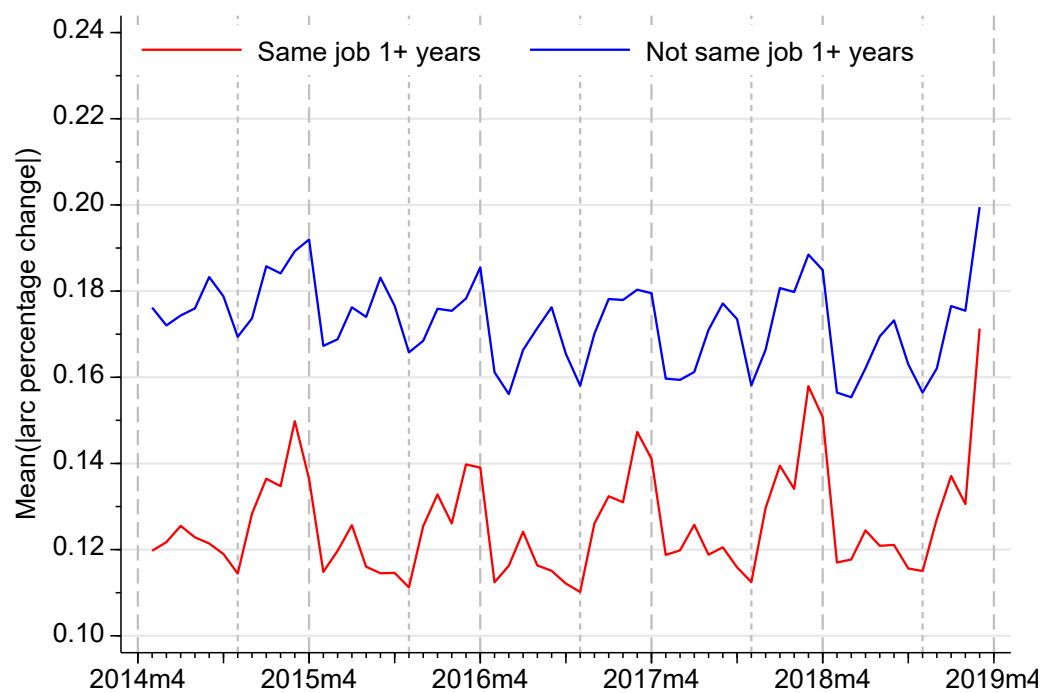
Source: HMRC PAYE data, HMRC file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t .

Figure B8. Monthly pay volatility, by whether has a temporary job contract



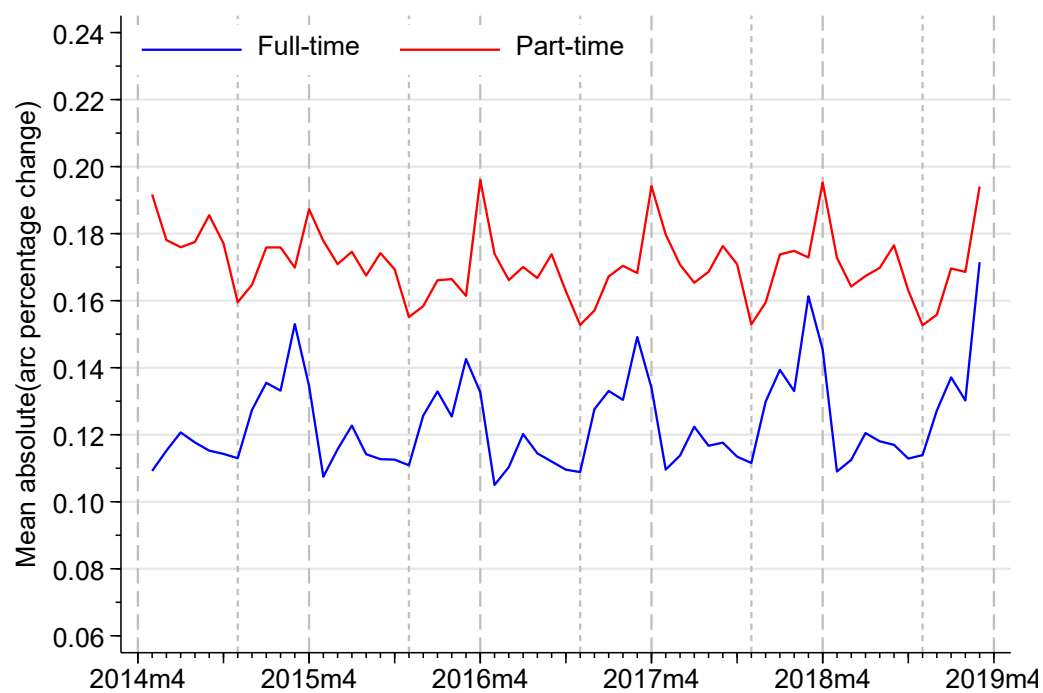
Source: HMRC PAYE data, Linked-ASHE file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t .

Figure B9. Monthly pay volatility, by whether in same job more than one year



Source: HMRC PAYE data, Linked-ASHE file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t .

Figure B10. Monthly pay volatility, by whether working full-time or part-time



Source: HMRC PAYE data, Linked-ASHE file. Notes. Calculations based on observations with positive earnings in month $t-1$ and t . Legend refers to age group (in years).

Appendix D. Twelve-month pay volatility, with subgroup breakdowns

Here we present estimates of volatility measured assessed using monthly earnings over a 12-month period, a multiperiod approach in contrast to the period-pair approach relied on so far. We are interested to see whether the longer assessment period makes a substantive difference to which groups are identified as having the greatest (or least) volatility.

We focus on the average CV (the cross-person average of the SD of monthly deviations from each individual's 12-month mean income normalized by that mean) and supplement it with a second measure, the fraction of individuals in a FY with four or more monthly deviations from their 12-month mean income that are at least 25% above or below that mean. Our calculations use a pooled sample of the five 12-month periods (FYs) covered by our dataset, and we use the Linked-ASHE file to exploit information about worker characteristics. Calculations are undertaken only for employees with positive earnings in each of the 12 months of the relevant FY. This sample selection is to better ensure that cross-group differentials calculated using the Linked-ASHE represent those in the HMRC file. (See Appendix Table E2 for some comparisons.)

Table C1 shows the estimates of our multiperiod volatility measures overall and separately for subgroups defined as earlier. (Characteristics refer to the situation in April each FY, the first month of the 12-month windows we are using.) For all persons in the Linked-ASHE file sample, the average CV is 0.171. It is difficult to assess whether this is a high or low number because there are few comparable benchmarks based on the same definition of 'income', window length, or sample selection. Angsten Clark and Otulana (2024) chart the income instability for five of their study participants, also showing the associated CV in each case. Hannagan and Morduch (2015) provide the following example to help assess CV estimates:

[C]onsider a hypothetical case in which a household's monthly income held steady at its average income during half of the year and then was 50 percent above average for the next three months and then 50 percent below average for the final three months.

That pattern roughly mirrors the pattern of spikes and dips in the USFD data [estimate = 38%], and it similarly generates a CV in the range here (35%). (Hannagan and Morduch, 2015, 8–9.)

Farrell et al. (2019, 14) report an average CV for gross family income of 0.48 using US bank account data. That estimate, and Hannagan and Morduch's hypothetical case estimate, are substantially larger than our estimate for individual earnings in the UK.

Perhaps easier to comprehend is the prevalence of 4+ large deviations measure, with an estimate of 0.122. That is, around one in eight persons experience four or more large deviations in their earnings around their 12-month average earnings per year. This is a large number.

<Table C1 near here>

The second panel of Table C1 shows that there are negligible differences between men and women according to the average CV, but women are more susceptible than men to 4+ large monthly earnings changes within a FY. The third panel shows that multiperiod earnings volatility is greatest for the youngest age group and declines with age. Again, differentials are magnified if we focus on the prevalence of 4+ large earnings changes.

The third panel provides a breakdown by decile group of 12-month-average earnings. (The subgroup sizes are not each 10% of the total because the Linked-ASHE file sample slightly under-represents the lowest earners.) Clearly, there is substantially greater multiperiod volatility at both the bottom and the top of the earnings distribution compared to the middle ranges. The averaged CV is much larger than the sample average for the bottom two groups and the top one, and below average for the other groups. The prevalence of large changes measure shows the same pattern, but the two lowest earnings groups stand out more as experiencing higher volatility. Almost a third of the bottom decile group experience 4+ large monthly earnings changes within a year, and a quarter of the next to bottom group.

Social care, hospitality, and leisure are the low-paid occupations with the highest levels of multiperiod volatility (as for period-pair volatility in Table 1). Looking at high-paid occupational groups, we see that all of them have average CV values above the sample average. On the other hand, only three groups (functional managers and directors; financial institution managers and directors; sales and marketing and related) have an above-sample average prevalence of 4+ large income changes. This is consistent with the other groups receiving large bonuses only once or maybe twice a year.

Individuals with monthly-type payslips and with weekly-type payslips have much the same volatility according to the average CV, which differs from what we found using a month-to-month measure – greater for the latter group. However, it is the weekly-type payslips employees who have an above-average prevalence of 4+ large monthly earnings changes within a year.

Table C1 also shows that that above average multiperiod volatility is experienced by workers with more than one job, on a temporary contract, paid hourly, working part-time, in the private sector, and not in the same job for at least a year. These differences correspond to

those revealed by the month-to-month volatility measures. Also, we find again that subgroup differentials are more pronounced in magnitude for the prevalence of 4+ large income changes measure than for the average CV.

The final panel of Table C1 shows multiperiod volatility differences for a selection of example person types ('profiles') and can be compared with the corresponding month-to-month estimates shown in Table D1 below. We see, for instance, that a lower bound reference point is provided by higher/further education professionals aged 45–49 who are not hourly paid. Less than 4% of this group experiences 4+ large pay deviations from their annual average. At the other extreme, the corresponding percentage for 20–24 year olds, paid hourly with weekly payslips, working in hospitality, is around 10 times larger, 42%.

<Table C1 near here>

Table C1. Multiperiod (twelve-month) earnings volatility, by subgroup

	Size (% of total)	Avg(CV)	%	Fraction with 4+ large deviations*	%
<i>Total (all working 12 months)</i>	100.00	0.171	100	0.122	100
<i>Sex</i>					
Female	48.16	0.167	98	0.139	113
Male	51.84	0.174	102	0.107	87
<i>Age (years)</i>					
20–24	11.75	0.205	120	0.218	178
25–29	16.27	0.180	105	0.137	112
30–34	16.17	0.180	105	0.129	105
35–39	13.39	0.174	102	0.116	95
40–44	11.62	0.163	95	0.099	81
45–49	11.86	0.155	90	0.087	71
50–54	10.75	0.148	87	0.085	69
55–59	8.19	0.146	85	0.088	72
<i>Decile group of pay (12-month average earnings)</i>					
D1 (lowest)	9.66	0.196	115	0.245	200
D2	9.42	0.190	111	0.238	194
D3	9.18	0.163	95	0.155	127
D4	9.19	0.141	83	0.098	80
D5	9.41	0.137	80	0.082	67
D6	9.83	0.134	78	0.069	57
D7	10.13	0.133	78	0.057	47
D8	10.43	0.135	79	0.056	46
D9	10.87	0.151	89	0.062	51
D10 (highest)	11.86	0.306	179	0.169	138
<i>Low-paid occupation</i>					
Non-low paying sectors	69.28	0.168	98	0.102	84
Agriculture	0.60	0.141	82	0.111	91
Food Processing	1.41	0.165	97	0.132	108
Textiles	0.22	0.144	84	0.088	72
Retail	8.20	0.184	107	0.171	140
Hospitality	4.38	0.201	117	0.231	189
Cleaning	2.61	0.178	104	0.177	145
Social care	2.79	0.202	118	0.231	189
Childcare	1.41	0.171	100	0.185	152
Leisure	0.80	0.199	116	0.213	174
Hairdressing	0.65	0.145	85	0.128	105
Office work	1.63	0.154	90	0.118	97
Non-food processing	1.62	0.161	94	0.108	88
Storage	2.56	0.164	96	0.123	100
Transport	1.84	0.138	81	0.090	73
<i>High-paid occupation</i>					
Not high paid	86.36	0.161	94	0.121	99
Functional managers & directors	2.65	0.301	176	0.168	138
Finan institution managers & directors	0.26	0.350	205	0.143	117
IT, telecoms	2.48	0.182	106	0.072	59
Business, research & admin professionals	2.31	0.208	121	0.091	75
Business, finance, etc., assoc professionals	1.98	0.234	137	0.121	99
Sales & marketing & related	3.96	0.238	139	0.168	138
<i>Pay slip type</i>					
All monthly	74.54	0.166	97	0.106	87
All weekly	22.71	0.181	106	0.161	132
Mixed or unknown	2.75	0.218	128	0.242	198
<i>Whether has 1+ jobs</i>					
One job	92.73	0.168	98	0.114	93

More than one job	7.27	0.213	124	0.227	186
<i>Contract type</i>					
Permanent	94.13	0.168	98	0.112	92
Temporary	5.87	0.224	131	0.285	233
<i>Hourly paid or not</i>					
Not hourly paid	60.97	0.169	99	0.100	81
Hourly paid	39.03	0.173	101	0.158	129
<i>Whether working full- or part-time</i>					
Part-time	24.53	0.189	111	0.214	175
Full-time	75.47	0.165	97	0.093	76
<i>Public/private sector job</i>					
Private	82.57	0.183	107	0.132	108
Public	17.43	0.114	66	0.076	62
<i>Same job more than one year</i>					
Same job 1+ years	80.77	0.166	97	0.108	89
Not same job	19.23	0.192	112	0.181	148
<i>Example person types</i>					
Not profile case	91.00	0.170	99	0.122	100
20–24, hourly-paid, weekly payslips, hospitality	0.72	0.281	164	0.418	342
35–39, !hourly-paid, monthly payslips, public sector	7.45	0.172	100	0.094	77
35–39, !hourly-paid, mthly payslips, private sector, sales & marketing	0.59	0.266	156	0.197	161
45–49, hourly-paid, weekly payslips, social care	0.13	0.174	102	0.164	134
45–49, !hourly-paid, monthly payslips, higher/further ed professionals	0.11	0.086	50	0.037	30

Source: HMRC PAYE data, Linked-ASHE file. Notes: Multiperiod volatility measures calculated using 12 months of FY data, pooling FYs. Estimates are weighted. All subgroup characteristics from the Linked-ASHE file refer to April of each survey year. High-paid subgroups correspond to SOC2010 groups 113x, 115x, 213x, 242x, 353x, and 354x, respectively. ‘%’ column shows subgroup volatility estimate as fraction of total volatility estimate. *: Fraction of individuals with 4+ monthly deviations at least 25% above or below their 12-month longitudinal mean earnings. ‘!’ means ‘not’.

Appendix D. Volatility profiles for selected person types and volatility regressions

Our subgroup breakdowns reveal that factors associated with high pay volatility levels are associated with being in the bottom tenth of the pay distribution or in a low-paid occupation, in a private sector job, paid by the hour and getting weekly pay packets, on a temporary contract, and being aged 20–24. The multiperiod volatility breakdowns point to most of the same factors as well.

These associations raise several questions. One is whether the factors have independent effects, e.g., what is the association between volatility and low-paid occupation if one adjusts for type of employment contract or whether hourly paid? To address the first question, we use three regressions – of the absolute arc percentage earnings change, the probability of having a large pay increase between month $t-1$ and t , and of the probability having a large pay decrease – pooling data from all 60 months (and including FY indicators). The regressors in each model are the characteristics used to define subgroups in Table 1, plus government region indicators. Estimated coefficients for each model represent marginal effects.

A second and related question is how volatility profiles look for individuals with specific combinations of characteristics. We consider this first.

Volatility profiles for selected person types

Summarizing volatility using $Mn(a_t)$, Figure D1 contrasts the experience of five ‘types’ with all other individuals (‘not profile case’), where types are constructed using assumptions about age, whether hourly paid, payslip type, and occupation and sector. The non-profile cases account for around 90% of the sample, and the next largest group (public sector workers) around 8%. The four other groups together account for around 2%. This is a further reminder that overall volatility depends not only on the volatility levels for specific subgroups but also subgroup size.

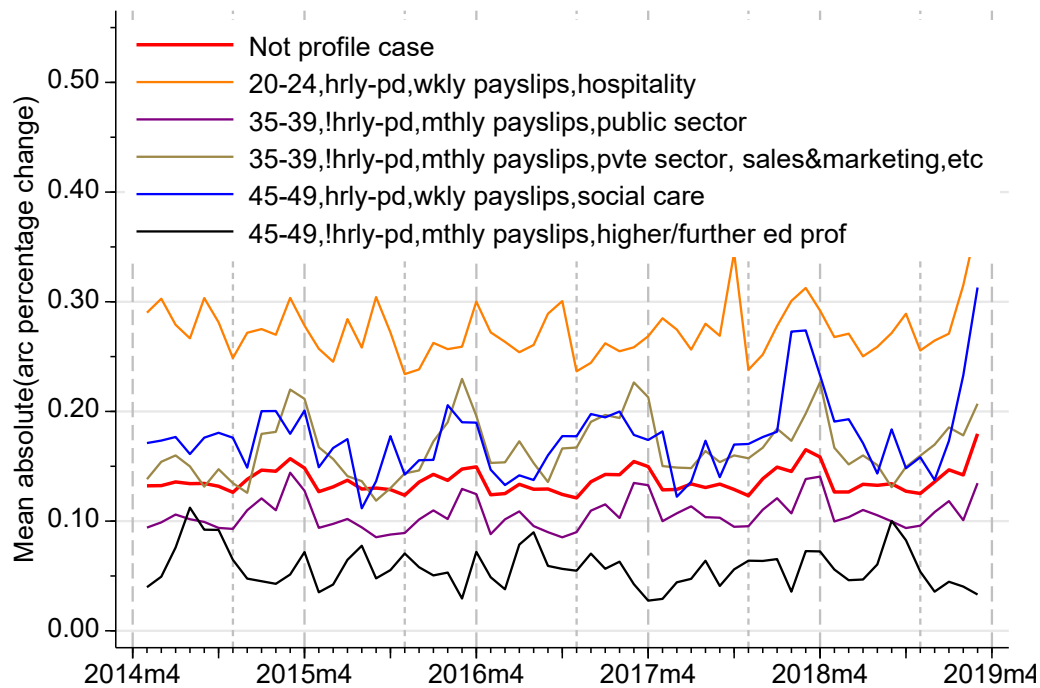
The non-profile cases have an average monthly percentage change of around 0.14 rising to a peak around April and trough in November. The volatility profile for those aged 35–39 working in the public sector is similar in shape but shifted down two to three percentage points (ppt). The cases with the lowest volatility are the salaried 45–49 year olds in higher/further education, 8 ppt lower than the non-profile cases at 0.06. The seasonality differs too: to the extent that there are peaks, they are in August, not April. In contrast, hourly-paid workers in hospitality aged 20–24 have substantially higher volatility than non-

profile cases, more than 10 ppt at 0.27. For them, there are peaks and troughs around April and November but also a mini-peak in October.

The remaining two cases illustrate how high volatility on average can arise in both low- and high-paid occupations. Average volatility for both the 35–39 year old in sales and marketing and the 45–49 year old social care worker is about 0.17, though observe that the amplitude of the seasonal cycle is larger for the latter.

<Figure D1 and Table D1 near here>

Figure D1. Monthly earnings volatility for selected subgroups of individuals



Source: HMRC PAYE data, Linked-ASHE file. Notes: Calculations use the Linked-ASHE file and ASHE weights, and refer to observations with positive earnings in months $t-1$ and t . ‘!’ in the legend means ‘not’.

Table D1. Volatility by profile type and index, all working at $t-1$ and t , pooled data

	Size (% of total)	Arc percentage change		Arc percentage change			
		Mean	%	SD	%	Variance	%
Total (all working at $t-1$ and t)	100.00	0.136	100	0.286	100	0.082	100
<i>Example person types</i>							
Not profile case	91.11	0.137	100	0.286	100	0.082	100
20–24, hourly-paid, weekly payslips, hospitality	0.74	0.274	201	0.442	155	0.195	238
35–39, !hourly-paid, monthly payslips, pub sector	7.33	0.107	78	0.258	90	0.066	81
35–39, !hourly-paid, mthly payslips, private sector, sales & marketing	0.59	0.167	123	0.344	120	0.118	144
45–49, hourly-paid, weekly payslips, social care	0.13	0.175	129	0.326	114	0.106	130
45–49, !hourly-paid, monthly payslips, higher/further ed professionals	0.10	0.057	42	0.174	61	0.030	37

Source: HMRC PAYE data, Linked-ASHE file. Notes: Person-month observations pooled over 60 months. Sample is observations with positive earnings in months $t-1$ and t . All estimates are weighted. ‘%’: subgroup estimate as % of corresponding All estimate (row 1). ‘!’ means ‘not’.

Regression estimates

The regression estimates are shown in Tables D2 and D3. The former table excludes earnings decile group membership indicators from the regressors, the latter includes them. In the regression analysis samples, the mean $|\text{arc percentage change}|$ is 0.14 (mean of dependent variable in Model 1), and the proportions with a monthly decrease or increase of at least 25% is around 0.08 (mean of the dependent variables in Models 2 and 3). Hence, and for brevity, we highlight only coefficient estimates (marginal effects) that are larger in absolute value than 0.02 for model 1 or larger in absolute value than 0.01 for models 2 and 3. Marginal effects are smaller for factors not mentioned in the commentary that follows.

Table D2 shows that in the mean $|\text{arc percentage change}|$ model (Model 1), there are marked effects are for selected months (April and especially March with higher volatility, and February with lower volatility). The absolute change declines with age up to around age 45. Absolute changes are larger for workers in low paid occupations such as hospitality, leisure, and especially social care, and for workers in high paid occupations such as functional managers and directors, sales and marketing, and especially financial institution managers. Absolute changes are also greater for workers with a weekly payslip type, with more than one job, working part-time, working in the public sector, not in the same job for more than a year, and especially workers with a temporary contract.

There are also marked effects apparent in the estimates of the probability of a large monthly decrease and a large increase. The chances of a large monthly *decrease* are greater in April, January, and February (~ 2 ppt), and August and September (~ 1 ppt), and some 2 to 3 ppt greater for workers who are aged 20–24 rather than older, or working in hospitality, social care, or leisure occupations. Among high-paid occupations, the marginal effect for being a financial institution manager/director is some +6 ppt and around +4 ppt for those in sales/marketing/etc. The chances of a large decrease are markedly greater for workers who have a weekly payslip type, are hourly paid, with more than one job, working part-time (~ 2 ppt), or in the private sector (~ 2 ppt) and higher still for workers on a temporary contract (~ 5 ppt).

For the probability of a large monthly *increase*, there are increases associated with April, July (~ 1 ppt), December (~ 2 ppt), and especially March (almost 6 ppt). Workers aged 20–24 have higher chances of a large increase than other age groups (~ 2 to 3 ppt), so too do those in hospitality, social care, or leisure (low-paid occupations). Among high paid occupations, higher chances of a large increases are for functional managers/directors,

business and finance associate professionals, sales and marketing (~ 4 ppt), and especially financial institution manager/directors (~ 7 ppt). The chances of a large increase are also markedly greater for those with more than one job (~ 8 ppt), working part-time, or in the private sector (~ 3 ppt) and higher still for workers on a temporary contract (~ 4 ppt).

The main consequence of adding earnings decile group membership to the regressors for each of the three models is for the estimates associated with low-paid and high-paid occupations, which is hardly surprising. See Table D3. (The marginal effects on other coefficients are much the same.) For example, in the probability of a large monthly *decrease* model, more low-paid occupations are associated with lower volatility (around 2 ppt for each) whereas those in the high-paid occupations have lower chances of a decrease (compared to Table D2). In the same model, however, being in one of the top earnings decile groups raises the chances of a decrease, by 18 ppt for the top tenth. In the probability of a large monthly *increase* model, most low-paid occupations are associated with lower chances and higher-paid occupations with higher chances (compared to Table D2). However, offsetting this, those in the bottom earnings tenth have substantially greater chances of a large increase than other tenths (ranging from 11 to 24 ppt depending on group). So, if someone is in a low-paid occupation and also in the bottom earning tenth, the combined effect implied by Table D3 is consistent with shown in Table D2.

The main conclusion to draw from the regressions is that the characteristics associated with higher volatility in the univariate analysis (charts and tables) remain important when adjusting for other characteristics.

Table D1. Volatility regressions

	arc % change (1)		Pr(large fall) (2)		Pr(large rise) (3)	
Month (ref. November)						
April	0.033 (0.001)	***	0.023 (0.001)	***	0.015 (0.001)	***
May	−0.004 (0.001)	***	0.000 (0.001)		0.008 (0.001)	***
June	−0.003 (0.001)	***	0.003 (0.001)	***	0.007 (0.001)	***
July	0.006 (0.001)	***	0.008 (0.001)	***	0.011 (0.001)	***
August	0.004 (0.001)	***	0.013 (0.001)	***	0.001 (0.001)	*
September	0.010 (0.001)	***	0.011 (0.001)	***	0.000 (0.001)	
October	0.007 (0.001)	***	−0.002 (0.001)	***	0.008 (0.001)	***
December	0.011 (0.001)	***	−0.001 (0.001)	**	0.020 (0.001)	***
January	0.018 (0.001)	***	0.026 (0.001)	***	0.002 (0.001)	***
February	0.020 (0.001)	***	0.023 (0.001)	***	−0.000 (0.001)	
March	0.040 (0.001)	***	−0.002 (0.001)	***	0.059 (0.001)	***
Financial year (ref. 2014/15)						
2015/16	−0.006 (0.000)	***	−0.002 (0.000)	***	−0.003 (0.000)	***
2016/17	−0.005 (0.000)	***	−0.001 (0.000)	***	−0.002 (0.000)	***
2017/18	−0.002 (0.000)	***	0.000 (0.000)		0.001 (0.000)	*
2018/19	−0.002 (0.000)	***	0.000 (0.000)		0.005 (0.000)	***
Man (ref. woman)	0.007 (0.000)	***	0.003 (0.000)	***	0.008 (0.000)	***
Age group (ref. 20–24)						
25–29	−0.028 (0.000)	***	−0.012 (0.000)	***	−0.016 (0.000)	***
30–34	−0.035 (0.000)	***	−0.015 (0.000)	***	−0.021 (0.000)	***
35–39	−0.043 (0.000)	***	−0.020 (0.000)	***	−0.025 (0.000)	***
40–44	−0.052 (0.000)	***	−0.025 (0.000)	***	−0.028 (0.000)	***
45–49	−0.057 (0.000)	***	−0.029 (0.000)	***	−0.031 (0.000)	***
50–54	−0.060 (0.001)	***	−0.030 (0.000)	***	−0.034 (0.000)	***
55–59	−0.057 (0.001)	***	−0.029 (0.000)	***	−0.035 (0.000)	***
Low paid occ. (ref. not low-paid occ.)						
Agriculture	0.001 (0.002)		−0.002 (0.001)		−0.000 (0.001)	
Food Processing	0.002 (0.001)	*	0.005 (0.001)	***	0.005 (0.001)	***
Textiles	−0.004 (0.003)		−0.002 (0.002)		−0.001 (0.002)	

Retail	-0.001 (0.000)	*	0.001 (0.000)	***	0.001 (0.000)	***
Hospitality	0.028 (0.001)	***	0.019 (0.001)	***	0.018 (0.001)	***
Cleaning	-0.002 (0.001)	***	0.002 (0.001)	**	-0.004 (0.001)	***
Social care	0.049 (0.001)	***	0.035 (0.001)	***	0.026 (0.001)	***
Childcare	0.008 (0.001)	***	0.007 (0.001)	***	0.007 (0.001)	***
Leisure	0.034 (0.001)	***	0.018 (0.001)	***	0.017 (0.001)	***
Hairdressing	-0.027 (0.002)	***	-0.017 (0.002)	***	-0.022 (0.002)	***
Office work	-0.010 (0.001)	***	-0.008 (0.001)	***	-0.007 (0.001)	***
Non-food processing	0.008 (0.001)	***	0.007 (0.001)	***	0.007 (0.001)	***
Storage	0.005 (0.001)	***	0.002 (0.001)	**	0.001 (0.001)	*
Transport	-0.005 (0.001)	***	-0.008 (0.001)	***	-0.009 (0.001)	***
High paid occ. (ref. not high-paid occ.)						
Functional managers & directors	0.037 (0.001)	***	0.019 (0.001)	***	0.022 (0.001)	***
Fin institution managers & directors	0.087 (0.002)	***	0.064 (0.002)	***	0.074 (0.002)	***
IT, telecoms	0.004 (0.001)	***	0.005 (0.001)	***	0.008 (0.001)	***
Business, res. & admin professionals	-0.005 (0.001)	***	-0.001 (0.001)		0.002 (0.001)	**
Business, fin., etc, assoc professionals	0.014 (0.001)	***	0.012 (0.001)	***	0.018 (0.001)	***
Sales & marketing & related	0.047 (0.001)	***	0.035 (0.001)	***	0.038 (0.001)	***
Payslip type (re. all monthly)						
All weekly	0.034 (0.000)	***	0.013 (0.000)	***	0.011 (0.000)	***
Mixed or unknown	0.019 (0.001)	***	0.008 (0.001)	***	0.015 (0.001)	***
Hourly paid (ref. not hourly paid)	0.017 (0.000)	***	0.010 (0.000)	***	0.009 (0.000)	***
More than 1 job in month (ref. 1 job)	0.079 (0.001)	***	0.031 (0.000)	***	0.076 (0.000)	***
Full-time (ref. part-time)	-0.046 (0.000)	***	-0.026 (0.000)	***	-0.029 (0.000)	***
Public sector (ref. private sector)	-0.027 (0.000)	***	-0.024 (0.000)	***	-0.027 (0.000)	***
Temporary contract (ref. permanent)	0.089 (0.000)	***	0.045 (0.000)	***	0.043 (0.000)	***
Not the same job 1+ years (ref. same job)	0.025 (0.000)	***	0.011 (0.000)	***	0.010 (0.000)	***
Region (ref. London)						
North East	-0.026 (0.001)	***	-0.012 (0.001)	***	-0.008 (0.001)	***
North West	-0.022 (0.000)	***	-0.009 (0.000)	***	-0.006 (0.000)	***
Yorkshire & Humberside	-0.027 (0.001)	***	-0.012 (0.000)	***	-0.009 (0.000)	***

East Midlands	-0.021 (0.001)	***	-0.009 (0.000)	***	-0.006 (0.000)	***
West Midlands	-0.020 (0.001)	***	-0.010 (0.000)	***	-0.007 (0.000)	***
South West	-0.022 (0.001)	***	-0.010 (0.000)	***	-0.007 (0.000)	***
East	-0.013 (0.000)	***	-0.006 (0.000)	***	-0.004 (0.000)	***
South East	-0.014 (0.000)	***	-0.006 (0.000)	***	-0.004 (0.000)	***
Wales	-0.024 (0.001)	***	-0.010 (0.001)	***	-0.007 (0.001)	***
Scotland	-0.023 (0.001)	***	-0.009 (0.000)	***	-0.005 (0.000)	***
Northern Ireland	0.050 (0.007)	***	0.035 (0.007)	***	0.023 (0.007)	***
Constant	0.200 (0.001)	***	0.099 (0.001)	***	0.096 (0.001)	***
R-squared	0.043		0.017		0.026	
Adj. R-squared	0.043		0.017		0.026	
No. person-months	6,530,583		6,467,224		6,467,224	

Source: HMRC PAYE data, Linked-ASHE file, individuals with positive earnings at $t-1$ and t . Notes. Models (2) and (3) estimated by OLS (linear probability models). Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Cells highlighted in yellow are those with absolute value of the coefficient estimate greater than 0.02 (model 1) or greater than 0.01 (models 2, 3) and statistically significant with $p < 0.001$. Means of dependent variables: 14% (model 1) and 7.7% (models 2 and 3).

Table D2. Volatility regressions, including earnings decile group as regressor

	arc % change (1)		Pr(large fall) (2)		Pr(large rise) (3)	
Month (ref. November)						
April	0.022 (0.000)	***	0.023 (0.001)	***	0.011 (0.001)	***
May	0.001 (0.000)	**	0.001 (0.000)	**	0.005 (0.000)	***
June	0.004 (0.000)	***	0.004 (0.000)	***	0.005 (0.000)	***
July	0.010 (0.000)	***	0.008 (0.000)	***	0.010 (0.000)	***
August	0.007 (0.000)	***	0.013 (0.001)	***	0.001 (0.000)	
September	0.008 (0.000)	***	0.011 (0.001)	***	−0.000 (0.000)	
October	0.004 (0.000)	***	−0.002 (0.001)	***	0.006 (0.000)	***
December	0.011 (0.000)	***	−0.001 (0.001)	**	0.019 (0.000)	***
January	0.020 (0.000)	***	0.025 (0.001)	***	0.003 (0.000)	***
February	0.019 (0.000)	***	0.022 (0.001)	***	0.000 (0.000)	
March	0.034 (0.000)	***	−0.002 (0.001)	***	0.056 (0.001)	***
Financial year (ref. 2014/15)						
2015/16	−0.003 (0.000)	***	−0.003 (0.000)	***	−0.001 (0.000)	***
2016/17	0.000 (0.000)		−0.002 (0.000)	***	0.002 (0.000)	***
2017/18	0.004 (0.000)	***	−0.000 (0.000)		0.004 (0.000)	***
2018/19	0.005 (0.000)	***	−0.001 (0.000)	*	0.009 (0.000)	***
Man (ref. woman)	0.007 (0.000)	***	−0.012 (0.000)	***	0.019 (0.000)	***
Age group (ref. 20–24)						
25–29	−0.017 (0.000)	***	−0.017 (0.000)	***	−0.008 (0.000)	***
30–34	−0.027 (0.000)	***	−0.027 (0.000)	***	−0.013 (0.000)	***
35–39	−0.036 (0.000)	***	−0.037 (0.000)	***	−0.015 (0.000)	***
40–44	−0.044 (0.000)	***	−0.045 (0.000)	***	−0.018 (0.000)	***
45–49	−0.047 (0.000)	***	−0.048 (0.000)	***	−0.021 (0.000)	***
50–54	−0.050 (0.000)	***	−0.048 (0.000)	***	−0.024 (0.000)	***
55–59	−0.050 (0.000)	***	−0.044 (0.000)	***	−0.028 (0.000)	***
Low paid occ. (ref. not low-paid occ.)						
Agriculture	0.004 (0.001)	***	0.028 (0.001)	***	−0.014 (0.001)	***
Food Processing	0.012 (0.001)	***	0.030 (0.001)	***	−0.006 (0.001)	***
Textiles	0.001		0.020	***	−0.014	***

	(0.002)		(0.002)		(0.002)	
Retail	-0.011	***	0.018	***	-0.027	***
	(0.000)		(0.000)		(0.000)	
Hospitality	0.006	***	0.037	***	-0.019	***
	(0.000)		(0.001)		(0.001)	
Cleaning	-0.026	***	0.023	***	-0.050	***
	(0.001)		(0.001)		(0.001)	
Social care	0.041	***	0.053	***	0.006	***
	(0.001)		(0.001)		(0.001)	
Childcare	-0.030	***	0.025	***	-0.049	***
	(0.001)		(0.001)		(0.001)	
Leisure	0.021	***	0.033	***	-0.005	***
	(0.001)		(0.001)		(0.001)	
Hairdressing	-0.052	***	-0.002		-0.063	***
	(0.001)		(0.002)		(0.002)	
Office work	-0.014	***	0.011	***	-0.026	***
	(0.001)		(0.001)		(0.001)	
Non-food processing	0.018	***	0.028	***	0.001	
	(0.001)		(0.001)		(0.001)	
Storage	0.014	***	0.026	***	-0.008	***
	(0.001)		(0.001)		(0.001)	
Transport	0.002	**	0.015	***	-0.019	***
	(0.001)		(0.001)		(0.001)	
High paid occ. (ref. not high-paid occ.)						
Functional managers & directors	-0.014	***	-0.055	***	0.022	***
	(0.001)		(0.001)		(0.001)	
Fin institution managers & directors	0.042	***	0.001		0.076	***
	(0.002)		(0.002)		(0.002)	
IT, telecoms	-0.019	***	-0.031	***	0.008	***
	(0.001)		(0.001)		(0.001)	
Business, res. & admin professionals	-0.020	***	-0.034	***	0.008	***
	(0.001)		(0.001)		(0.001)	
Business, fin., etc, assoc professionals	0.009	***	-0.006	***	0.022	***
	(0.001)		(0.001)		(0.001)	
Sales & marketing & related	0.031	***	0.009	***	0.040	***
	(0.001)		(0.001)		(0.001)	
Earnings decile group (ref. D1)						
D2	-0.133	***	-0.001	*	-0.124	***
	(0.000)		(0.000)		(0.000)	
D3	-0.179	***	-0.000		-0.189	***
	(0.000)		(0.000)		(0.000)	
D4	-0.199	***	0.002	***	-0.216	***
	(0.000)		(0.001)		(0.001)	
D5	-0.203	***	0.011	***	-0.227	***
	(0.000)		(0.001)		(0.001)	
D6	-0.203	***	0.020	***	-0.234	***
	(0.000)		(0.001)		(0.001)	
D7	-0.200	***	0.031	***	-0.238	***
	(0.000)		(0.001)		(0.001)	
D8	-0.195	***	0.045	***	-0.239	***
	(0.001)		(0.001)		(0.001)	
D9	-0.183	***	0.069	***	-0.239	***
	(0.001)		(0.001)		(0.001)	
D10	-0.091	***	0.179	***	-0.230	***

	(0.001)		(0.001)		(0.001)	
Payslip type (re. all monthly)						
All weekly	0.033	***	0.017	***	0.004	***
	(0.000)		(0.000)		(0.000)	
Mixed or unknown	0.025	***	0.011	***	0.015	***
	(0.001)		(0.001)		(0.001)	
Hourly paid (ref. not hourly paid)	0.016	***	0.016	***	0.004	***
	(0.000)		(0.000)		(0.000)	
More than 1 job in month (ref. 1 job)	0.089	***	0.026	***	0.090	***
	(0.000)		(0.000)		(0.000)	
Full-time (ref. part-time)	0.026	***	-0.046	***	0.056	***
	(0.000)		(0.000)		(0.000)	
Public sector (ref. private sector)	-0.023	***	-0.022	***	-0.021	***
	(0.000)		(0.000)		(0.000)	
Temporary contract (ref. permanent)	0.055	***	0.046	***	0.033	***
	(0.000)		(0.000)		(0.000)	
Not the same job 1+ years (ref. same job)	0.015	***	0.016	***	0.005	***
	(0.000)		(0.000)		(0.000)	
Region (ref. London)						
North East	-0.012	***	0.004	***	-0.014	***
	(0.001)		(0.001)		(0.001)	
North West	-0.009	***	0.007	***	-0.012	***
	(0.000)		(0.000)		(0.000)	
Yorkshire & Humberside	-0.012	***	0.005	***	-0.015	***
	(0.000)		(0.000)		(0.000)	
East Midlands	-0.010	***	0.006	***	-0.013	***
	(0.000)		(0.000)		(0.000)	
West Midlands	-0.010	***	0.005	***	-0.013	***
	(0.000)		(0.000)		(0.000)	
South West	-0.009	***	0.006	***	-0.013	***
	(0.000)		(0.000)		(0.000)	
East	-0.007	***	0.003	***	-0.009	***
	(0.000)		(0.000)		(0.000)	
South East	-0.006	***	0.003	***	-0.008	***
	(0.000)		(0.000)		(0.000)	
Wales	-0.010	***	0.007	***	-0.013	***
	(0.001)		(0.001)		(0.001)	
Scotland	-0.008	***	0.004	***	-0.007	***
	(0.000)		(0.000)		(0.000)	
Northern Ireland	0.034	***	0.057	***	0.007	
	(0.006)		(0.007)		(0.007)	
Constant	0.276	***	0.083	***	0.228	***
	(0.001)		(0.001)		(0.001)	
R-squared	0.098		0.040		0.066	
Adjusted R-squared	0.098		0.040		0.066	
No. person-months	6,409,159		6,409,159		6,409,159	

Source: HMRC PAYE data, Linked-ASHE file, individuals with positive earnings at $t-1$ and t . Notes. Models (2) and (3) estimated by OLS (linear probability models). Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Cells highlighted in yellow are those with absolute value of the coefficient estimate greater than 0.02 (model 1) or greater than 0.01 (models 2, 3) and statistically significant with $p < 0.001$. Means of dependent variables: 14% (model 1) and 7.7% (models 2 and 3).

Appendix E. Comparisons of volatility distributions and volatility breakdowns: HMRC versus Linked-ASHE file estimates

Here we report our checks of how the volatility distributions and volatility breakdowns by subgroups that we derive using the Linked-ASHE file (reported in the main text) correspond with the corresponding distributions and breakdowns that we derive using the HMRC file.


Our checks are restricted to individuals working in months $t-1$ and t (month-pair analysis) or working all 12 months in a FY (multi-period analysis) because we know that the Linked-ASHE file is more non-representative of the HMRC file ‘parent’ if we include months with zero employment earnings (for the reasons given in the main text).

For the comparisons of subgroup breakdowns, we are of course restricted to defining subgroups using the limited set of characteristics that are present in both the HMRC file and the Linked-ASHE file: sex, age, earnings, payslip type, and whether has more than one job.


In what follows, we present corresponding estimates (summarised in a chart or table) in pairs, first showing the HMRC file estimate and then the Linked-ASHE file estimate. For brevity, there is no commentary, but we think it is easily apparent that, given the sample restriction to workers only, there is a close correspondence between pairs of estimates. On this basis, we claim that Linked-ASHE file breakdowns by subgroups defined using characteristics that are only available in that file are also reliable. This claim cannot be tested!

Figure E1. The distribution of arc percentage changes in monthly earnings (fractions per band)

(a) HMRC file




(b) Linked-ASHE file




Source: HMRC PAYE data, HMRC and Linked-ASHE files, 60 months pooled.

Figure E2. Empirical cumulative distribution functions for arc percentage changes in monthly earnings

(a) HMRC file



(b) Linked-ASHE file



Source: HMRC PAYE data, HMRC and Linked-ASHE files, 60 months pooled. Notes. ECDFs estimated using 100 equally-sized bins.

Table E1. Summary statistics for distributions on arc percentage changes in real earnings

(a) Real earnings: HMRC file

Sample	%	Mean	SD	p_{50}	p_{10}	p_{25}	p_{75}	p_{90}	Skewness	Kurtosis	Fraction with zero
All	100.00	0.003	0.461	0.000	-0.192	-0.012	0.022	0.205	-0.029	14.481	0.222
All working at $t-1$ and t	76.48	0.002	0.316	-0.002	-0.209	-0.036	0.042	0.219	-0.031	14.784	0.042

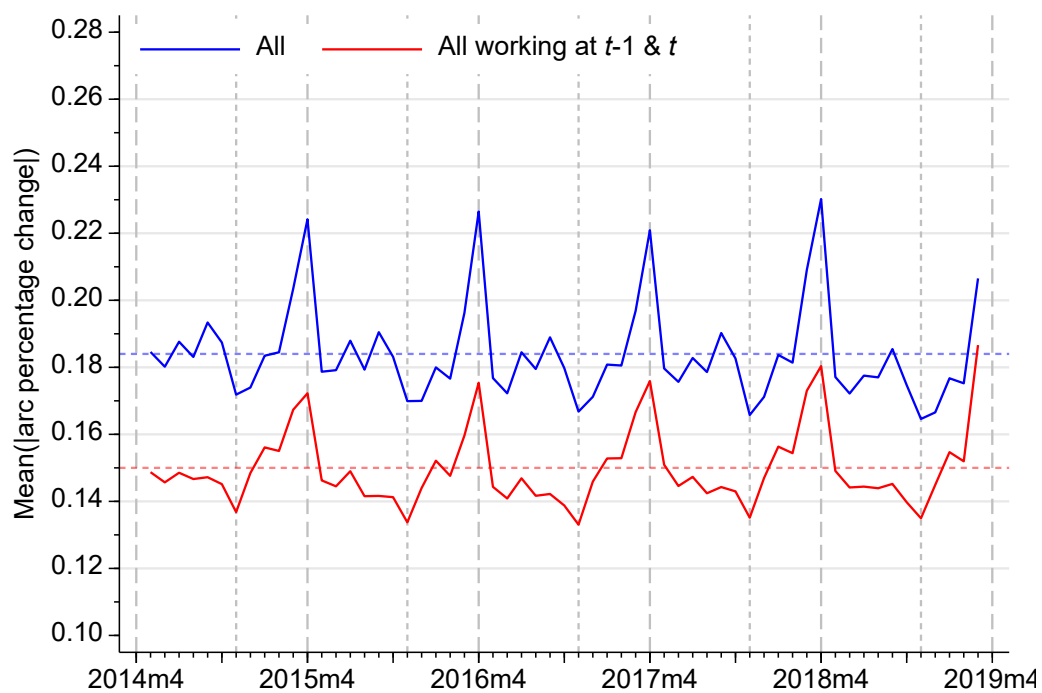
(b) Linked-ASHE file

Sample	%	Mean	SD	p_{50}	p_{10}	p_{25}	p_{75}	p_{90}	Skewness	Kurtosis	Fraction with zero
All	100.00	-0.011	0.411	-0.001	-0.215	-0.036	0.037	0.206	-0.593	16.444	0.074
All working at $t-1$ and t	92.90	-0.001	0.286	-0.002	-0.195	-0.039	0.040	0.199	-0.226	16.329	0.043


Source: HMRC PAYE data, HMRC and Linked-ASHE files, 60 months pooled.

Figure E2. Monthly earnings volatility, all and working at $t-1$ and t

(a) HMRC file




(b) Linked-ASHE file



Note: the horizontal dotted lines show the 60-month pooled average of the series.

Figure E3. Monthly earnings volatility by sex

(a) HMRC file



(b) Linked-ASHE file





Figure E4. Fractions with monthly earnings increase and decrease of at least 25%

(a) HMRC file



(b) Linked-ASHE file





Figure E5. Monthly earnings volatility, by age group

(a) HMRC file



(b) Linked-ASHE file





Figure E6. Monthly earnings volatility, by monthly/weekly payslip type

(a) HMRC file



(b) Linked-ASHE file





Figure E7. Monthly earnings volatility, by whether has more than one job

(a) HMRC file



(b) Linked-ASHE file






Figure E8. Monthly earnings volatility, by decile group of earnings at $t-1$

(a) HMRC file



(b) Linked-ASHE file



Note: decile groups defined using HMRC file (chart a) and Linked-ASHE file (chart b)

Table E2. Period-pair volatility breakdowns: HMRC file versus Linked-ASHE file comparisons, sample is all working in months $t-1$ and t

(a) HMRC file

	Size (% of total)	Arc percentage change		Arc percentage change			
		Mean	%	SD	%	Variance	%
All (all working at $t-1$ and t)	100.00	0.150	100	0.316	100	0.100	100
<i>Sex</i>							
Female	49.76	0.149	100	0.316	100	0.100	100
Male	50.24	0.150	100	0.315	100	0.100	100
<i>Age (years)</i>							
20–24	11.78	0.223	148	0.406	129	0.165	165
25–29	14.19	0.167	112	0.334	106	0.111	111
30–34	13.88	0.154	103	0.320	101	0.102	102
35–39	12.79	0.143	95	0.307	97	0.094	94
40–44	12.29	0.133	89	0.293	93	0.086	86
45–49	12.97	0.126	84	0.282	89	0.080	80
50–54	12.39	0.122	81	0.276	87	0.076	76
55–59	9.69	0.125	83	0.283	90	0.080	80
<i>Pay frequency type</i>							
All monthly	73.48	0.124	82	0.279	88	0.078	78
All weekly	23.42	0.214	143	0.390	123	0.152	152
Mixed or unknown	3.10	0.284	189	0.462	146	0.214	214
<i>Whether has 1+ jobs</i>							
One job	91.79	0.139	93	0.300	95	0.090	90
More than one job	8.21	0.269	179	0.452	143	0.204	204
<i>Decile group of pay at $t-1$</i>							
D1 (lowest)	10.00	0.371	248	0.587	186	0.345	345
D2	10.00	0.143	96	0.295	93	0.087	87
D3	10.00	0.147	98	0.281	89	0.079	79
D4	10.00	0.121	81	0.243	77	0.059	59
D5	10.00	0.107	71	0.222	70	0.049	49
D6	10.00	0.104	69	0.216	68	0.047	47
D7	10.00	0.101	67	0.213	67	0.045	45
D8	10.00	0.099	66	0.212	67	0.045	45
D9	10.00	0.106	71	0.230	73	0.053	53
D10 (highest)	10.00	0.196	131	0.392	124	0.154	154

(b) Linked-ASHE file

	Size (% of total)	Arc percentage change		Arc percentage change			
		Mean	%	SD	%	Variance	%
All (all working at $t-1$ and t)	100.00	0.136	100	0.286	100	0.082	100
<i>Sex</i>							
Female	48.52	0.137	101	0.290	101	0.084	102
Male	51.48	0.134	99	0.282	99	0.080	97

<i>Age (years)</i>							
20–24	12.53	0.186	137	0.348	122	0.121	148
25–29	16.58	0.146	107	0.296	103	0.088	107
30–34	16.18	0.137	101	0.288	101	0.083	101
35–39	13.28	0.129	95	0.279	98	0.078	95
40–44	11.39	0.123	90	0.270	94	0.073	89
45–49	11.52	0.117	86	0.259	91	0.067	82
50–54	10.45	0.114	83	0.255	89	0.065	79
55–59	8.08	0.116	85	0.262	91	0.068	83
<i>Pay frequency type</i>							
All monthly	74.82	0.118	87	0.265	93	0.070	86
All weekly	22.27	0.179	132	0.330	116	0.109	133
Mixed or unknown	2.91	0.248	182	0.411	144	0.169	206
<i>Whether has 1+ jobs</i>							
One job	92.24	0.127	94	0.274	96	0.075	92
More than one job	7.76	0.233	171	0.401	140	0.161	196
<i>Decile group of pay at $t-1$, defined using Linked-ASHE file sample</i>							
D1 (lowest)	10.00	0.283	208	0.494	173	0.244	297
D2	10.00	0.155	114	0.290	101	0.084	102
D3	10.00	0.128	94	0.248	87	0.061	75
D4	10.00	0.107	79	0.219	77	0.048	58
D5	10.00	0.100	74	0.207	73	0.043	52
D6	10.00	0.096	71	0.203	71	0.041	50
D7	10.00	0.095	70	0.202	71	0.041	50
D8	10.00	0.095	70	0.205	72	0.042	51
D9	10.00	0.104	76	0.225	79	0.051	62
D10 (highest)	10.00	0.192	141	0.381	133	0.145	177
<i>Decile group of pay at $t-1$, defined using HMRC file sample</i>							
D1 (lowest)	8.33	0.315	232	0.525	184	0.276	336
D2	7.79	0.152	112	0.293	102	0.086	105
D3	9.16	0.141	103	0.266	93	0.071	86
D4	9.84	0.114	84	0.228	80	0.052	64
D5	10.18	0.102	75	0.211	74	0.045	54
D6	10.71	0.098	72	0.205	72	0.042	51
D7	11.14	0.096	70	0.203	71	0.041	50
D8	11.29	0.094	69	0.203	71	0.041	50
D9	11.31	0.103	76	0.223	78	0.050	61
D10 (highest)	10.25	0.190	140	0.378	132	0.143	174

Source: HMRC PAYE data, HMRC and Linked-ASHE files. Statistics based on those working in months $t-1$ and t . Pooled data for 60 months. Linked-ASHE calculations are weighted. Breakdowns are for characteristics present in both files.

Table E3. Twelve-month volatility breakdowns (all working all 12-months in FY): HMRC file versus Linked-ASHE file

(a) HMRC file

	Size (% of total)	Avge(CV)	%	Fraction with 4+ large deviations*	%
Total (all working 12 months in FY)	100.00	0.175	100	0.132	100
<i>Sex</i>					
Female	49.44	0.169	97	0.143	108
Male	50.56	0.180	103	0.122	92
<i>Age (years)</i>					
20–24	10.19	0.217	124	0.244	185
25–29	13.71	0.188	107	0.155	117
30–34	13.79	0.187	107	0.148	112
35–39	12.96	0.179	102	0.130	99
40–44	12.72	0.169	97	0.112	85
45–49	13.59	0.159	91	0.097	74
50–54	13.02	0.152	87	0.093	70
55–59	10.03	0.149	85	0.095	72
<i>Decile group of pay (12-month average earnings)</i>					
D1 (lowest)	10.00	0.182	104	0.227	172
D2	10.00	0.183	104	0.227	172
D3	10.00	0.178	102	0.187	142
D4	10.00	0.150	86	0.115	87
D5	10.00	0.141	81	0.090	68
D6	10.00	0.138	79	0.078	59
D7	10.00	0.136	78	0.065	49
D8	10.00	0.137	79	0.060	46
D9	10.00	0.158	90	0.071	54
D10 (highest)	10.00	0.344	196	0.203	154
<i>Pay slip type</i>					
All monthly	73.75	0.167	95	0.111	84
All weekly	23.23	0.192	110	0.184	139
Mixed or unknown	3.02	0.229	131	0.262	198
<i>Whether has 1+ jobs</i>					
One job	91.92	0.171	97	0.123	93
More than one job	8.08	0.223	127	0.243	184

(b) Linked-ASHE file

	Size (% of total)	Avge(CV)	%	Fraction with 4+ large deviations*	%
Total (all working 12 months)	100.00	0.171	100	0.122	100
<i>Sex</i>					
Female	48.16	0.167	98	0.139	113
Male	51.84	0.174	102	0.107	87
<i>Age (years)</i>					
20–24	11.75	0.205	120	0.218	178
25–29	16.27	0.180	105	0.137	112
30–34	16.17	0.180	105	0.129	105
35–39	13.39	0.174	102	0.116	95
40–44	11.62	0.163	95	0.099	81
45–49	11.86	0.155	90	0.087	71
50–54	10.75	0.148	87	0.085	69
55–59	8.19	0.146	85	0.088	72
<i>Decile group of pay (12-month average earnings)</i>					

D1 (lowest)	9.66	0.196	115	0.245	200
D2	9.42	0.190	111	0.238	194
D3	9.18	0.163	95	0.155	127
D4	9.19	0.141	83	0.098	80
D5	9.41	0.137	80	0.082	67
D6	9.83	0.134	78	0.069	57
D7	10.13	0.133	78	0.057	47
D8	10.43	0.135	79	0.056	46
D9	10.87	0.151	89	0.062	51
D10 (highest)	11.86	0.306	179	0.169	138
<i>Pay slip type</i>					
All monthly	74.54	0.166	97	0.106	87
All weekly	22.71	0.181	106	0.161	132
Mixed or unknown	2.75	0.218	128	0.242	198
<i>Whether has 1+ jobs</i>					
One job	92.73	0.168	98	0.114	93
More than one job	7.27	0.213	124	0.227	186

Source: HMRC PAYE data, HMRC and Linked-ASHE files. Statistics based on those working all 12 months in a FY. Pooled data for 5 FYs. Linked-ASHE calculations are weighted. Breakdowns are for characteristics present in both files. *: Fraction with 4+ monthly deviations at least 25% of 12-month mean