

Are the Self-Employed More Exposed to Income Volatility than Employees? A Panel Analysis of Employment Status in the UK.

Adwoa Owusu
Professional Economist BSc and Apprenticeship Level 6
School of Economics
University of Kent, 2026

Executive summary

This study investigates the extent to which self-employed individuals are more exposed to income volatility than employees in the United Kingdom (UK). The analysis uses panel data from the UK Household Longitudinal Study covering over 50,000 individuals. Income volatility is measured as the absolute arc-percentage change in annual net labour income, with the main explanatory variable being self-employment status. Several econometric approaches are taken, including a pooled ordinary least squares and fixed effects model. A two-way fixed effects model with individual and time fixed effects controls for unobserved heterogeneity and time-invariant differences across approximately 4,000 individuals who transition between employment statuses during this period. Results show a consistent and statistically significant relationship between self-employment and income volatility and are robust to alternative methods. Specifically, self-employed individuals face almost 31 percentage points higher income volatility compared to a period of wage employment. This indicates that the volatility difference is not solely driven by time-invariant characteristics, though selection based on time-varying factors cannot be ruled out. Additional models explore heterogeneity across occupation, education and gender. The increase in income volatility associated with self-employment is more pronounced for intermediate- and high-skilled occupations. In contrast, there is no robust evidence that this difference varies by education or gender. In a period of expansion in self-employment, this raises questions about the economic risks faced by self-employed individuals. The findings highlight the importance of distinguishing between employment statuses when assessing income volatility and suggest that understanding the drivers of such differences remains an important area for future research.

AI Statement

I acknowledge the use of generative AI in code development and drafting in this paper. However, the work reported remains my own.

Acknowledgements

I would like to acknowledge my colleagues at the Bank of England for their unwavering support, encouragement and guidance throughout my apprenticeship. Additionally, I would like to thank my line manager of almost four years for her dedication to my professional and academic development and for motivating me to continuously improve. Thank you also, to the University of Kent for the academic support over the duration of this programme. I could not have completed this journey without the invaluable love and constant support of my friends and family; my siblings and niblings (nieces and nephews) for their inspiritment and my parents, who instilled lifelong principles in me to fear God, learn hard and work hard. Finally, I owe my deepest gratitude to my Lord and Saviour Jesus Christ. 1 John 4:10.

1. Introduction

1.1 Background

Self-employment has become an increasingly prominent feature of the United Kingdom (UK) labour market. From representing 9% of the workforce in the early 1970s, the share of self-employed individuals rose to 14% by 2019, when self-employment peaked at almost 5 million individuals (ONS, 2026a; ONS, 2026b). This expansion has been characterised by a growth in solo self-employment, platform-based work and the gig economy, enabling new forms of labour market participation characterised by flexible hours and independent contractual work (ONS, 2019; Giupponi and Xu, 2020).

These structural changes have sparked interest in how economic risk is distributed across the labour market, especially as non-standard work frequently comes at the cost of income security, shifting risk from firms to individual workers (Freeland and Prassl, 2017). A key dimension of this risk is income volatility. Unlike income levels, which capture long-run differences, income volatility reflects short-term instability and uncertainty, with direct implications for individual welfare and financial resilience.

For the purposes of this paper, self-employment refers to individuals whose main economic activity is running a business or professional practise, working for themselves, freelancing, or operating as a subcontractor. Income volatility refers to the degree to which individual income fluctuates over time. The study focuses on individual- rather than household-level income, to isolate the direct effect of employment status on personal economic security, abstracting from the risk-sharing and income smoothing that occurs within households. Establishing whether self-employment entails greater income volatility is essential for understanding its welfare consequences and how well existing institutions protect financially vulnerable workers.

1.2 Research question and hypotheses

The central research question is:

To what extent are self-employed individuals more exposed to income volatility than employees in the UK?

Whilst the existing literature has examined income volatility and household earnings dynamics, there is comparatively less evidence on individual-level income volatility by

employment status, particularly when accounting for employment transitions. This paper contributes to the literature by focusing on post-tax income, exploiting panel variation to identify within-person changes associated with switching into self-employment, and exploring heterogeneity across occupation, education and gender.

The analysis uses panel data from the UK Household Longitudinal Study (UKHLS), from 2009 to 2019. Income volatility is measured as the absolute value of the arc-percentage change in annual net labour income.

The paper investigates two hypotheses:

Hypothesis 1:

- **H1:** Self-employed individuals experience higher income volatility than employees.
- **H1a:** Within the same individual, income volatility is higher in a period of self-employment than in wage employment.

Hypothesis 2:

- **H2:** The income volatility premium associated with self-employment differs across occupation, education and gender. Specifically, the difference is:
 - **H2a:** Larger for low-skilled individuals than for high-skilled individuals.
 - **H2b:** Larger at lower levels of education relative to higher levels of education.
 - **H2c:** Larger for females relative to males.

The remainder of this paper is structured as follows: Section 2 reviews the theoretical and empirical literature on self-employment and income volatility. Section 3 describes the data, followed by the methodology in Section 4. Section 5 presents the results, while Section 6 discusses the policy implications of these results. Section 7 considers limitations and finally, Section 8 concludes the study, with suggestions for future research.

2 Literature Review

2.1 Theoretical Framework

Two prominent labour market arrangements are wage employment (working as an employee) and self-employment. These differ not only in how earnings are generated but also in the degree of income stability, access to institutional protection, and the mechanisms through

which economic shocks are absorbed. In wage employment, earnings are determined by organisational pay structures and employment contracts, accompanied by predictable pay schedules and statutory protections. By contrast, self-employed individuals rely more directly on market demand, client relationships, and business performance, facing weaker institutional insulation from revenue fluctuations. If exposure to economic shocks differ systematically across these arrangements, employment status may act as a structural determinant of income volatility rather than merely correlating with it.

Income risk and consumption smoothing

The welfare implications of income volatility are grounded in intertemporal consumption theory. The Permanent Income Hypothesis (PIH) (Friedman, 1957) and the Life-Cycle Hypothesis (LCH) (Modigliani and Brumberg, 2005) both model how individuals smooth consumption over time but differ in emphasis. The PIH holds that consumption is based on expected long-run or “permanent” income, so temporary fluctuations should have little effect on spending. The LCH emphasises predictable changes in income over the lifetime, arguing that individuals borrow early in life, save during peak earning years, and dissave in retirement to maintain stable consumption. Both imply that if individuals can freely borrow and save, temporary income shocks should have limited welfare consequences.

In practise, however, real economies have incomplete markets, where credit constraints and income uncertainty limit the ability to smooth consumption (Carroll, 1992; Deaton, 1992). The precautionary savings model (Carroll 1997) addresses this by predicting that individuals facing greater income risk accumulate “buffer-stock savings” as a form of self-insurance. However, where individuals face liquidity constraints, precautionary savings may be insufficient to fully offset income shocks. In such cases, income volatility, particularly when persistent or large, can translate into consumption instability and welfare loss. This is particularly relevant for individuals whose earnings are directly exposed to market fluctuations, such as the self-employed, and motivates the central focus of this dissertation.

Labour market structure and risk allocation

In wage employment, organisational pay structures often dampen the immediate transferral of economic shocks to workers’ earnings. The implicit contract theory (Gordon, 1974; Azariadis 1975) argues that risk-averse workers may prefer stable wages, as firms, being better able to diversify risk, can absorb short-term fluctuations in demand or productivity. As a result, wages are smoothed over time, with profits rather than wages adjusting to temporary shocks.

Employees in the UK additionally benefit from statutory redundancy protections, sick pay entitlements, and access to social insurance, all of which further insulate earnings from shocks.

Unlike employees, self-employed individuals lack such institutional smoothing. Income is more directly tied to business performance, customer demand, pricing conditions and payment timing (Knight, 1921; Parker, 2009). Firm-level risk pooling, where shocks are absorbed across many workers, is mostly absent. Consequently, labour market theory implies that the self-employed should experience higher income volatility than employees (Knight, 1921).

However, this contrast should not be overstated. Employees remain exposed to income risk through wage variation (Low, Meghir and Pistaferri, 2010), technological change, and job loss (Jacobson, LaLonde and Sullivan, 1993). Moreover, some self-employed individuals – especially those with stable client bases or incorporated structures – may experience relatively predictable income streams (Parker, 2009), but when earnings depend on business revenue rather than contractual wages, exposure to market shocks increases (Knight, 1921).

Accordingly, if markets are incomplete and if income determination differs across employment status, theory predicts systematic differences in volatility between employees and the self-employed, therefore motivating H1. This perspective connects to broader accounts of labour market restructuring which argue that organisational changes such as contracting, gig-work and the expansion of self-employment, have progressively shifted economic risk from employers towards workers (Friedman, 2014; Freeland and Prassl, 2017).

Occupational choice and selection into self-employment

Observed volatility differences may not reflect institutional risk reallocation alone but may also reflect who enters self-employment. Under the expected utility theory (von Neumann, Morgenstern, and Rubinstein, 1944), individuals evaluate income risk based on expected utility, with choices depending on their degree of risk aversion. The compensating differentials theory (Smith, 1776) similarly suggests that riskier occupations offer higher pay or non-pecuniary benefits, such as autonomy and flexibility, to attract workers. If self-employment carries greater income risk, rational individuals select into it only if they expect to be compensated financially or psychologically. Parker (2009) applies the expected utility theory, where concave, linear and convex functions represent risk-averse, risk-neutral and

risk-loving individuals respectively (Figure 1). Risk-averse individuals prefer stable earnings and are more likely to choose wage employment, while more risk-tolerant individuals select into self-employment. This implies that the self-employed are typically more willing to bear income fluctuations, although real-world choices are also shaped by non-risk-related factors.




Figure 1: Utility functions (Parker, 2009)

Sorting models reinforce this point. The Roy (1951) model explains occupational choice through comparative advantage across sectors, while Lucas (1978) focuses on differences in managerial ability, with higher-ability individuals selecting into self-employment. While both models are useful for understanding selection into self-employment, neither map cleanly onto contemporary UK self-employment, much of which involves solo or low-income work rather than the entrepreneurship central to Lucas. The Roy model also disregards institutional constraints and non-pecuniary motives. Taken together, these models suggest that higher observed volatility may reflect the preferences and characteristics of those who select into self-employment rather than a causal effect of employment status. This is explored in H1a.

The push-pull model challenges the voluntary narrative by distinguishing between individuals who are ‘pulled’ into self-employment by opportunities, autonomy, and other non-pecuniary benefits, and those who are ‘pushed’ into self-employment out of necessity due to limited employment prospects, unemployment or weak labour demand (Storey, 1991; Amit and Muller, 1995; Dawson, Henley, and Latreille, 2009). Push-motivated entrants are likely to face systematically higher income instability and fewer resources to buffer shocks.

Empirically, failing to account for selection may lead to attributing differences in income volatility to employment status when they may instead reflect unobserved heterogeneity among individuals who choose self-employment.

Human capital and heterogeneity

The human capital theory (HCT) (Schultz, 1961; Becker 1964) suggests that differences in education, skills and experience contribute to income dispersions as human capital raises productivity and stabilises earnings. Applied to self-employment, however, this is less clear. While highly educated or skilled self-employed individuals may have greater earning capacity, they may also face more variable income from project-based work and fluctuating demand. As a result, income volatility is unlikely to be uniform in the self-employed population and may vary by skill level.

Differences in human capital, alongside demographic characteristics like gender, may further complicate this picture, as labour market segmentation and differential access to networks can produce distinct volatility profiles within the self-employed population. These considerations underpin H2 and suggest heterogeneity within the self-employed, carrying important implications for the empirical analysis.

2.2 Empirical framework

Measurement of Income Volatility

The empirical literature uses several approaches to measure income volatility, each with distinct trade-offs. Early work by Gottschalk et al. (1994) separated income into permanent and transitory components. However, this method assumes that income is the sum of distinct and typically uncorrelated components. In practice, temporary shocks may persist or influence long-run earnings (Meghir and Pistaferri, 2004). Additionally, the use of log earnings typically excludes zero-income observations, potentially understating the true extent of income volatility for those more prone to earning gaps (Ziliak et al., 2011; Celik et al., 2012).

Subsequent research has developed summary measures of volatility based on year-to-year income changes. These include the standard deviation of the change in log earnings (Celik et al., 2012; Shin and Solon, 2011), and three-year log rolling standard deviations (Dynan et al., 2007, 2012). A widely used alternative is the standard deviation of the arc percentage change in income (Dynan et al., 2007; Dahl et al., 2011; Ziliak et al., 2011; Cappellari and Jenkins, 2014; Avram et al., 2021), which divides the change in income by the average of two consecutive years. This measure reduces sensitivity to outliers and, critically, allows the inclusion of zero-income observations – an important advantage when analysing employment and income volatility.

Despite these developments, the literature remains divided over whether volatility should include zero earnings. Cappellari and Jenkins (2014) show this distinction is not trivial; when zero earnings are included, measured volatility reflects not only wage fluctuations but also transitions into and out of employment. Summary measures also have limitations because unlike permanent-transitory decompositions, they do not distinguish persistent structural shocks from short-term fluctuations. Consequently, studies that exclude labour market exits may underestimate economic instability, while those including them may conflate earnings volatility with employment risk. These methodological choices therefore have implications for both estimated volatility levels and comparisons across employment statuses.

Cross-country income volatility trends

Empirical evidence on income volatility trends in the UK present a more nuanced picture than the well-documented rise in the US. Using British Household Panel Survey data, Cappellari and Jenkins (2014) found that earnings volatility among workers with positive earnings in consecutive years remained broadly stable over the 1990s and 2000s. However, when including labour market transitions, volatility declined over the same period, reflecting stronger labour market attachment rather than changes in within-job earnings volatility.

More recent evidence based on UKHLS data similarly documents a modest decline in individual earnings and disposable household income volatility (Avram et al., 2021) but contrasts with the US literature, which documents substantial increases in earnings instability since the 1970s (Gottschalk and Moffitt, 2009; Dynan et al, 2012). Differences may reflect the UK's labour market institutions, welfare systems and employment protection legislation, which differ in structure and operation.

While the existing literature is concentrated in the US and the UK, a smaller body of research examines income volatility in other countries. European panel data suggest that earnings instability is shaped by labour market institutions, with stronger employment protections and higher union density associated with lower volatility (Sologon and O'Donoghue, 2014). In contrast, research from developing economies find that higher volatility is driven by exposure to economic shocks in contexts where formal insurance mechanisms are absent (Dercon, 2004). This is compounded in economies with large informal markets, where workers lack access to social insurance against unemployment and income loss, leaving them disproportionately exposed to income volatility (Perry et al., 2007).

However, cross-country comparisons are limited by differences in data availability, measurement approaches, and sample restrictions, particularly for countries where reliable longitudinal income data are scarce. Consequently, much of the current understanding of income volatility is derived from Western labour markets, highlighting the importance of institutional context when interpreting results.

Self-employment and income volatility

Empirical evidence consistently suggests that income during self-employment spells fluctuate more than earnings in wage employment (Hamilton, 2000; Parker, 2009). Using panel data, Audoly (2024) demonstrates that earnings shocks during self-employment spells are substantially larger than those in wage employment, with a significantly higher probability of large negative income shocks. Studies using UK panel data similarly report greater self-employment income volatility relative to the wage equivalent (Avram et al., 2021). Research examining entry into self-employment finds that transitions from employee jobs are often associated with sustained earnings declines, even where individuals report higher job satisfaction and autonomy (Clark and Oswald, 1994; Blanchflower and Oswald, 1998; Hundley, 2001; Benz and Frey, 2008; Cribb and Xu, 2020).

The literature also emphasises that volatility within the self-employed is not uniform. Necessity-driven entrants tend to experience higher levels of income instability than opportunity-driven entrepreneurs (Glavin, Filipovic, and van der Maas, 2019; Moulton and Scott, 2016), and lower-income self-employed households face higher odds of experiencing volatile income streams (Auguste, Roll and Despard, 2022). A key methodological limitation is that many studies rely on cross-sectional comparisons of income distributions rather than within-person changes over time. This creates an identification problem as it is difficult to determine true income volatility within the self-employed population. Self-employment income is also more prone to measurement error due to under-reporting, which may inflate volatility estimates (Brandolini et al., 2004; Cabral, Gemmell and Alinaghi, 2020). These limitations mean that existing estimates of volatility differences between employment statuses should be interpreted with caution and underscore the importance of using panel methods in the empirical analysis.

The role of institutions

Institutions play a central role in shaping observed volatility. Avram et al. (2021) highlight the stabilising role of the UK tax-benefit system, which accounted for most of the reduction between gross and net income volatility, a finding which supports the role of institutions discussed in the theoretical framework.

Additionally, wage employment typically provides more stable earnings structures, while self-employment income fluctuates with business performance and market demand (Evans and Leighton, 1990; Hamilton, 2000; Audoly, 2024). Self-employed workers lack access to institutional cushions, particularly those below the National Insurance (NI) lower earnings limit or those whose variable income complicates benefit eligibility, leaving them more exposed to demand and revenue shocks (Parker, 2009; Moulton and Scott, 2016). This asymmetry is cited as a key reason why self-employment income tends to be more volatile (Parker, 2009; OECD, 2019).

Some self-employed households may adopt alternative risk-management strategies, smoothing consumption by drawing on savings, working longer hours or relying on household income diversification (Parker et al., 2005; McKnight and Rucci., 2020). These mechanisms may partially offset higher income volatility, though their effectiveness depends on household wealth and access to financial resources.

2.3 Gaps in literature and contributions

Despite the growing literature, several gaps remain. Much of the UK evidence relies on studies that either examine volatility at the household level (obscuring individual-level employment status differences), or do not directly compare within-person income changes across employment statuses over time. In addition, relatively little empirical work examines how volatility differs by occupations, education levels, and genders. This dissertation contributes to the literature by examining individual-level income volatility differences between employees and the self-employed in the UK, using panel data methods that account for selection and the role of occupational, educational, and gender differences.

3 Data

3.1 Dataset and sample

The analysis uses the UKHLS (University of Essex, 2026), a large-scale nationally representative longitudinal survey following approximately 40,000 households and up to 100,000 individuals annually since 2009. The survey collects detailed individual-level information on labour market status, employment type, income and a range of demographic and health characteristics, making it suitable for constructing individual income volatility measures and comparing these across employment statuses over time.

The analysis focuses on Waves 1 to 11, corresponding to the period 2009 to 2019. This window provides a sufficiently long panel to construct individual-level volatility, while avoiding the confounding effects of the COVID-19 pandemic. From Wave 12 onwards, the introduction of large-scale government support schemes - including the Coronavirus Job Retention Scheme (furlough) and the Self-Employment Income Support Scheme - temporarily smoothed income for both employment statuses. These interventions introduced dynamics which were unrepresentative of structural differences across employment statuses (Blundell et al., 2020). Excluding post-2019 waves avoids the risk of conflating policy-induced income stabilisation with underlying employment-type differences in income risk.

The sample is restricted to individuals aged 16 to 64, who report being in either self-employment or wage employment as their main economic activity at the time of interview. This age group corresponds to the standard working-age population in the UK and excludes individuals approaching or past the state pension age, whose labour market attachment likely differs qualitatively. All other economic statuses, including unemployed, retired, and full-time students are also excluded as they are not reflective of permanent labour market attachment.

After applying these restrictions and removing observations with missing values on the income variable and key covariates, the final sample consists of 51,368 individuals. The sample is unbalanced, as not all individuals are observed in every wave, which is addressed in the construction of the volatility measure described below.

3.2 Dependent variable

3.2.1 Income variable

The income variable used to construct the volatility measure is annual net labour income, defined as earnings from employment or self-employment after tax and NI contributions, excluding non-labour income. Income is adjusted for inflation using the Consumer Prices Index (Housing) to ensure values are comparable in real terms.

Net rather than gross income is used for two reasons. First, net income better reflects the disposable resources available to individuals for consumption smoothing. Second, asymmetric tax treatment of employees and the self-employed in the UK, including differences in NI contributions and deductible business expenses, mean that gross income measures may combine differences in tax exposure with income volatility. Net income standardises for these differences to the extent they are captured in reported take-home pay, though limitations of self-reported income data are acknowledged in Section 7.

3.2.2 Measuring income volatility

Income volatility is measured as the absolute value of the arc-percentage change in annual net labour income. The arc-percentage change between periods t and $t-1$ for individual i is defined as:

$$\Delta y_{it} = 100 \times \frac{Y_{it} - Y_{it-1}}{[(Y_{it} + Y_{it-1})/2]}$$

Where Y_{it} denotes individual i 's net annual labour income in year t . The arc-percentage change is preferred to log differences for two reasons. First, it accommodates zero and negative income values, which can arise for the self-employed when business losses offset other earnings and which log-transformations exclude by construction. Second, by dividing by the average of two periods, this reduces sensitivity to extreme values in either direction, which are prevalent among the self-employed (Dahl et al., 2011).

Much of the existing empirical methods use the standard deviation of this arc-percentage change. While informative, this approach collapses the panel into a single summary statistic per individual, therefore disregarding the within-person variation. In contrast, the absolute value retains individual-year observations. Taking the absolute value focuses on the magnitude of income volatility regardless of direction, allowing the distribution of within-person volatility to be analysed. Figure 2 highlights this distinction.

The standard deviation compresses variation, whereas the distribution of the absolute value has a more pronounced right skew. Both distributions show the self-employed exhibit a wider spread and heavier tails than employees. These differences are descriptive and unconditional, and may reflect observable differences between groups. The following section investigates whether this relationship persists after controlling for such factors.

Therefore, following Key, Prager and Burns (2017), income volatility (V_i) is then defined as the absolute value of these arc-percentage changes across all consecutive year pairs for which the individual has valid income observations:




Figure 2: Distribution of income volatility by employment type

$$V_i = |\Delta y_{it}|$$

A minimum of two consecutive income observations is required for the absolute value to be computed, however, individuals with fewer than three observations are excluded from the sample to allow multiple observations per individual. The panel is unbalanced as individuals

contribute different numbers of change observations depending on their survey participation, but this is preferable to restricting the sample to a fully balanced panel, which would introduce selection bias by retaining only those with uninterrupted survey participation across all waves.

3.3 Independent variables

Table 1 provides a summary of all variables included in the analysis.

Table 1: Variable descriptions

Variable	Description	Type	Summary Statistics				
			Mean	Standard deviation	Min	Max	Obs
Income volatility	Absolute value of the arc-percentage change in net labour income.	Continuous	23.867	34.551	0	200	176,785
Self-employed	Binary indicator which equals 1 if the individual reports self-employment as their main employment status.	Binary	0.132	0.339	0	1	250,434
Occupation	Occupational group from low- to high-skilled (1 = Sales & Operatives; 2 = Administrative & skilled; 3 = Managerial & professional)	Categorical	2.221	0.806	1	3	245,674
Education	Highest qualification gained (1 = no qualification; 2 = GCSE/equivalent; 3 = A-level/equivalent; 4 = Higher education or Degree).	Ordinal	4.330	1.482	1	6	247,511
Potential experience Potential experience ²	Calculated as Age – Schooling – 6 (Mincer, 1974)	Continuous	23.500	12.324	0	54	249,716
Female	Binary indicator equals 1 if the individual gender is female.	Binary	0.509	0.500	0	1	250,424
Health	Self-assessed general health status rated on a 1-5 scale (1 = poor; 5 = excellent).	Ordinal	3.692	0.973	1	5	136,092
Region	Nine Government Office Regions of England, plus Scotland, Wales and Northern Ireland.	Categorical (12 dummies)	6.601	3.127	1	12	250,300
Ethnicity	Ethnic group recoded into five broad categories: White British; Asian/Asian British; Black/Black British; Mixed; Other.	Categorical (dummies)	1.486	1.087	1	5	244,135

4. Methodology

4.1 Model development

The empirical analysis is split into three stages, each addressing part of the research question. The first stage estimates a pooled ordinary least squares (OLS) regression to understand whether self-employed individuals experience greater income volatility than employees on average. It establishes a descriptive baseline comparable to the existing UK literature and provides the raw employment status volatility gap. The equation takes this form:

$$V_i = \alpha + \beta_1 X_i + \beta_2 X_i + \beta_3 X_i + \dots + \beta_8 X_i + \varepsilon_i$$

However, collapsing the panel to one observation creates an identification problem by conflating within-person income volatility with persistent between-person income volatility. Furthermore, OLS cannot control for time-invariant unobserved heterogeneity, such as risk tolerance and ability. The OLS estimate of β_1 is therefore expected to be upward biased relative to the true within-person effect.

The second stage addresses the identification limitations by preserving the full panel structure and estimating a two-way fixed effects (TWFE) model. The main specification is:

$$V_i = \alpha_i + \beta_1 \text{self-employed}_{it} + \gamma X_{it} + \delta_t + \varepsilon_{it}$$

where α_i are individual fixed effects, *self-employed*_{it} is the primary independent variable, *X*_{it} contains time varying controls, and δ_t are time fixed effects. Individual fixed effects absorb all time-invariant unobserved heterogeneity, including risk tolerance, ability and demographic characteristics, which simultaneously drive selection into self-employment and income volatility. Time fixed effects account for macroeconomic shocks such as inflation and policy changes.

The coefficient β_1 identifies the within-person effect of transitions into and out of self-employment on income volatility. This is estimated from the 3,961 individuals who switch employment status at some point during the panel.

The third stage tests whether the self-employment premium, defined as the difference in income volatility between the self-employed and employees, varies with occupation, education, and gender, as predicted by the HCT (Schultz, 1961; Becker, 1964). Education is largely time-invariant within the 2009-2019 panel and is therefore absorbed by individual fixed effects. A standard approach for testing how a time-invariant moderator influences a treatment effect is the use of interactions (Wooldridge, 2019). The TWFE from stage two is




Figure 3: Variable correlation matrix

therefore estimated separately in three models, each with a different interaction term. The category ‘other qualification’ is excluded from the education analysis on the grounds that it is insufficiently specific to place in the education categories. Comparing the β_1 coefficients across the three models tests whether the self-employment volatility premium differs, as H2 predicts.

4.2 Diagnostic tests

Several diagnostic tests were conducted to verify the assumptions underlying each model.

4.2.1 OLS diagnostics

Starting with diagnostics for the pooled OLS specification, the Breusch-Pagan test strongly rejects homoskedasticity ($p < 0.01$) confirming that robust standard errors are necessary.

A correlation matrix is used to check for potential multicollinearity (Figure 3). The high correlation between potential experience and its squared term reflects the functional relationship, so is not a significant concern.

In addition, variance inflation factors are generally low with a mean of 5.37, indicating there is no problematic multicollinearity among the main explanatory variables.

Tests for residual normality were assessed by examining a histogram, which showed a right-skew. However, given the large number of observations ($N > 75,000$), this is unlikely to affect inference due to the central limit theorem.

The Ramsey RESET test rejects the null of no omitted variables ($p < 0.01$), suggesting possible model misspecification in the pooled OLS framework. While not uncommon in large datasets, this result highlights the limitations of this specification and motivates the use of alternative modelling approaches.

4.2.2 Panel diagnostics

To assess whether a panel model is required, the Breusch-Pagan Lagrangian Multiplier test was conducted. The test rejects the null of no panel effects ($p < 0.01$), indicating the presence of unobserved individual heterogeneity and that panel estimators are preferred over pooled OLS.

The Hausman test was used to determine the appropriate panel estimator. The test ($\chi^2 = 413, p = 0.000$) formally rejects the random effects estimator, confirming that individual-level fixed characteristics are correlated with self-employment status and the fixed effects model is the appropriate specification.

The Wooldridge test detects first-order autocorrelation ($p < 0.01$), supporting the use of clustered standard errors at the individual level. This is further supported by the modified Wald test for groupwise heteroskedasticity in the fixed effects model, which indicates that there is heteroskedasticity across individuals ($p < 0.01$). Clustered standard errors are therefore used in all models to address heteroskedasticity and serial autocorrelation across observations from the same person.

5. Results

This section presents the results from the three-stage empirical strategy, exploring how employment status influences income volatility. The analysis establishes the self-employment volatility gap, alongside within-person identification, removing the influence of time-invariant individual characteristics.

Table 2: Income volatility and self-employment

	(1)	(2)	(3)	(4)	(5)
	Pooled OLS	TWFE (baseline)	TWFE + SE × Occupation	TWFE + SE × Education	TWFE + SE × Female
<i>DV: V_i</i>					
Self-employed (β_1)	46.380*** (0.745)	30.725*** (1.895)	23.744*** (3.640)	20.939** (9.948)	28.326*** (2.151)
Occupation (ref: Low-skilled)					
Intermediate	-1.860*** (0.413)	-2.133* (1.143)	-3.192*** (1.028)	—	—
High-skilled	-3.223*** (0.415)	-2.205* (1.195)	-3.461*** (1.113)	—	—
SE × Intermediate	—	—	8.045* (4.427)	—	—
SE × High-skilled	—	—	7.585* (4.135)	—	—
Education (ref: No qualification)					
GCSE	-1.061 (0.780)	5.636 (7.564)	—	3.801 (7.429)	—
A-level	-0.930 (0.796)	0.017 (7.507)	—	-1.697 (7.423)	—
Higher/Degree	-1.411* (0.804)	6.338 (8.171)	—	5.104 (8.081)	—
SE × GCSE	—	—	—	13.493 (10.913)	—
SE × A-level	—	—	—	10.804 (10.564)	—
SE × Higher/Degree	—	—	—	8.114 (10.275)	—
Gender (ref: Male)					
Female	1.714*** (0.282)	Absorbed	Absorbed	Absorbed	Absorbed
SE × Female	—	—	—	—	4.158 (3.932)
Potential experience	-0.828*** (0.051)	0.219 (0.620)	-0.348 (0.461)	0.206 (0.613)	-0.375 (0.453)
Potential experience ²	0.015*** (0.001)	0.013** (0.005)	0.012*** (0.005)	0.014*** (0.005)	0.013*** (0.005)
Health	Yes	Yes	Yes	Yes	Yes
Ethnicity	Yes	Absorbed	Absorbed	Absorbed	Absorbed
Region FE	Yes	Yes	Yes	Yes	Yes
Individual FE	No	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes
Standard error	Robust	Clustered	Clustered	Clustered	Clustered
Mean (DV)	23.404	23.404	23.526	23.452	23.579
N	74,704	74,704	80,568	75,357	81,272
R ²	0.199	0.023 (within)	0.023 (within)	0.023 (within)	0.022 (within)

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Region dummies included throughout; coefficients omitted for brevity as all are non-significant in TWFE specifications. In all specifications, “—” indicates variables excluded; “Absorbed” indicates variables absorbed by individual fixed effects.

Extensions to the baseline explore how the self-employment premium varies across occupation, education and gender, directly engaging with the discussions set out in the

theoretical framework. Across these specifications, the self-employment coefficient is consistently positive and significant, suggesting that self-employed individuals experience higher income volatility than employees. Finally, robustness checks are reported to assess the stability of the main findings across alternative specifications and samples. The main results are summarised in Table 2 above.

5.1 Stage 1: The Raw Self-employment Volatility Gap

The starting point is the pooled OLS regression reported in column (1) of Table 2. The coefficient on self-employment (46.380, $p < 0.01$) indicates that the self-employed face income volatility approximately 46 percentage points larger than employees, subject to controls. With a dependent variable mean of approximately 23 percentage points, the self-employment premium represents an increase in volatility of nearly twice the average level. This is consistent with the UK evidence from Audoly (2024) and Avram et al. (2021) and provides strong initial support for H1.

The control variables broadly align with theory and literature. The occupational gradient is negative and statistically significant. Relative to low-skilled workers, those in intermediate occupations experience around 1.9 percentage points less income volatility, while high-skilled workers face a reduction of around 3.2 percentage points (both $p < 0.01$). This aligns with the prediction of the HCT that higher-skilled individuals experience more stable earnings. Importantly, occupation explains a larger share of the income volatility than education. While education coefficients are negative, they are small and mostly insignificant, with only higher-degree level education showing weak significance (-1.411, $p < 0.10$). This suggests that the nature of the work itself, rather than the credentials leading to it, most directly determine volatility.

Potential experience follows the expected life-cycle pattern with a negative linear term (-0.828, $p < 0.01$) and positive quadratic term (0.015, $p < 0.01$), indicating that income volatility declines through mid-career, before rising slightly in later working life as employment attachment weakens. The relationship between health and income volatility is non-linear. Relative to individuals reporting poor health, those reporting 'fair' and 'good' health show no significant differences, while individuals reporting 'very good' and 'excellent' health show higher volatility (up to 3 percentage points), significant at least at the 5% level.

Due to its limitations, the OLS estimate of 46 percentage points is best understood as an upper bound on the causal effect. The comparison with the TWFE results allows the selection bias to be isolated.

5.2 Stage 2: Within-Person Identification

Column (2) of Table 2 introduces individual fixed effects. The coefficient on self-employment falls to 30.725 ($p < 0.01$), identified from the 3,961 individuals who transition between employment statuses during the observation period. This result directly supports H1_a: within the same individual, periods of self-employment are associated with income volatility approximately 30.7 percentage points higher than periods of wage employment, even after controlling for personal characteristics. The reduction from the OLS estimate of 46.4 percentage points to the TWFE estimate of 30.7 percentage points implies a selection bias of approximately 15.7 percentage points. This gap may represent stable characteristics like risk tolerance and unobserved ability, which simultaneously predict self-employment entry and income volatility. However, the possibility that selection bias varies across individuals and across time cannot be disregarded.

The occupation coefficients remain negative (Intermediate: -2.133; High-skilled: -2.205; both $p < 0.10$) and directionally consistent with the OLS, though reduced in significance as individual fixed effects absorb much of the stable between-person occupational sorting. Education produces positive but entirely non-significant coefficients, with large standard errors. Although there is within-person variation in educational attainment, these results are expected as the changes are confined to a small group of individuals who gain qualifications during the observation period. The positive signs cannot be interpreted as evidence that acquiring higher qualifications increases income volatility. Regional coefficients are similarly non-significant – individuals rarely change region within the panel, leaving insufficient variation to identify regional effects.

The linear potential experience term is non-significant in the TWFE, in contrast to its significance in the OLS. This is an expected consequence of the fixed effects design as during a ten-year period, within-person variation in experience is limited, reducing its explanatory power. However, the squared term remains positive and significant (0.013, $p < 0.05$), indicating a mild increase in volatility in later careers. Health is typically stable for individuals over short-to-medium timeframes, and the within-person variation that does exist is insufficient to produce precise estimates, hence its non-significance in the TWFE.

Gender and ethnicity are time-invariant and therefore absorbed by individual fixed effects. Wave fixed effects show a consistent declining pattern from 2011 onwards, reflecting the post-financial crisis stabilisation of earnings documented by Cappellari and Jenkins (2014).

The low within- R^2 (0.023) is expected in a fixed effects model and reflects the limited share of within-person changes explained by self-employment and time-varying controls.

5.3 Stage 3: Heterogeneity by Occupation, Education, and Gender

The mostly non-significant occupation and education coefficients in the baseline TWFE reflect the low within-person variation available to separately identify these effects alongside individual fixed effects. To test whether the self-employment premium varies across these dimensions, columns (3) – (5) introduce interaction specifications building on the baseline TWFE. These directly address Hypothesis 2.

Column (3) introduces interactions between self-employment and occupation. The main self-employment coefficient (23.744, $p < 0.01$) now represents the premium for low-skilled individuals, implying that this group experiences 23.7 percentage points higher income volatility when self-employed than when employed. The occupation main effects remain negative and significant, indicating that among employees, higher-skilled workers experience lower income volatility. The interaction terms for intermediate (8.045, $p < 0.10$) and high-skilled (7.585, $p < 0.10$) occupations are positive and marginally significant, implying that the within-person self-employment premium is larger in those groups relative to low-skilled individuals, with total effects rising to 31.8 and 31.3 percentage points respectively. This contradicts H2_a: while self-employment increases income volatility across all occupations, the effect is amplified among intermediate and high-skilled workers. Conditional on individual fixed effects, intermediate- and high-skilled occupations are associated with lower baseline volatility, however, entry into self-employment is associated with the greatest increase in income volatility among these workers.

The effects of education interactions can be seen in column (4). The self-employment coefficient suggests that switching into self-employment increases income volatility by around 20.9 percentage points for individuals with no qualifications, significant at the 5% level. The interaction terms for GCSE (13.494), A-level (10.804) and Higher/Degree-level education (8.114) are positive but statistically insignificant, indicating no robust evidence that the self-employment premium differs systematically by education level. An alternative

approach using stratified education samples yields larger and statistically significant coefficients, but these estimates are not directly comparable due to differences in sample composition (see Appendix). As such, the results are inconclusive and H2_b is not supported.

Finally, column (5) examines gender heterogeneity. The baseline self-employment coefficient implies men experience around 28 percentage points higher income volatility when self-employed compared to when they are employees. The female interaction coefficient (4.158) is positive but statistically insignificant, providing no robust evidence that the increase in income volatility associated with self-employment differs between men and women. The contrast with the positive female coefficient in the OLS specification possibly suggests that the observed gender difference reflects time-invariant characteristics rather than differences in how self-employment affects income volatility across gender. Accordingly, H2_c is not supported.

5.4 Robustness Checks

Table 3 presents the robustness of the main findings across seven specifications. In each case, the self-employment coefficient remains positive, large and statistically significant, supporting H1 and confirming that the core result is insensitive to variation in methodology.

Table 3: Robustness checks

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
	SD arc-%	3yr rolling SD	Dominant SE	≥5 waves	Winsorised 1 st /99 th	Balanced panel	arc-% gross income
	<i>OLS</i>	<i>OLS</i>	<i>OLS</i>	<i>TWFE</i>	<i>TWFE</i>	<i>TWFE</i>	<i>TWFE</i>
Self-employed (β_1)	52.061***	2.310***	41.761***	30.717***	29.782***	34.329***	33.078***
	(0.807)	(0.082)	(0.734)	(2.041)	(1.822)	(4.376)	(1.995)
Potential experience	-0.481***	-0.003	-0.802***	0.238	0.206	0.210	0.695
	(0.060)	(0.004)	(0.052)	(0.660)	(0.609)	(1.277)	(0.662)
Potential experience ²	0.009***	0.000	0.015***	0.012**	0.013**	0.001	0.013***
	(0.001)	(0.000)	(0.001)	(0.005)	(0.005)	(0.011)	(0.005)
Health	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Individual FE	No	No	No	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	98,830	47,209	74,704	60,445	74,704	14,410	74,704
R ²	0.273	0.178	0.170	0.024	0.023	0.030	0.024

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Three OLS checks assess sensitivity to the dependent variable and employment classification. Column (i) replicates the established literature, using the standard deviation of arc-percentage changes in net income ($\beta_1 = 52.061$, $p < 0.01$). Column (ii) employs a three-year rolling

standard deviation as the dependent variable to capture more persistent volatility patterns. This result is directionally consistent with the main result. Column (iii) classifies self-employment by an individual's dominant status across the panel ($\beta_1 = 41.761, p < 0.01$), confirming results are not sensitive to whether employment status is measured at each point in time or is the same across the panel.

Four TWFE checks address sample composition and outlier sensitivity. Column (iv) restricts to individuals with at least five consecutive wave pairs ($\beta_1 = 30.717, p < 0.01$). This is virtually identical to the main estimate, confirming thin panel histories are not material. Column (v) winsorises the absolute arc-percentage change in net income at the 1st and 99th percentiles ($\beta_1 = 28.782, p < 0.01$), confirming outliers do not drive the result. Column (vi) restricts to a balanced panel observed across all eleven waves, yielding $\beta_1 = 34.329 (p < 0.01)$. The larger estimate in the balanced sample reflects the fact that individuals continuously observed for a decade include a disproportionate share of long-term self-employed workers with more persistently volatile income. Even so, the self-employment volatility is large and persistent across both samples (Figure 4).




Figure 4: Income volatility by employment status (2009 - 2019)

Note: 95% confidence intervals.

Finally, column (vii) uses the absolute arc-percentage change in gross income as the dependent variable, allowing comparison with pre-tax income.

6. Policy implications

The finding that switching into self-employment is associated with a substantial increase in income volatility has several policy implications. Most directly, it suggests that existing institutional frameworks may be less well suited to individuals with irregular earnings. Systems such as means-tested benefits that assess eligibility based on current income, may provide inconsistent support for those whose incomes fluctuate over time, potentially leaving gaps when financial vulnerability is greatest.

The heterogeneity in the results indicate that exposure to income risk is uneven across the self-employed population. This points to the potential value of more targeted forms of support that help individuals manage income fluctuations, as uniform policy approaches may not fully address variation by occupation and skill level.

More broadly, increased income volatility may influence participation in self-employment, potentially deterring entry into self-employment and encouraging exit among those less able to absorb financial risk. This could have consequences for the entrepreneurial activity that contributes to innovation and job creation. Policies which better accommodate income variability, such as aligning benefit assessments with fluctuating earnings or structuring tax and support systems to smooth income over time, may help reduce these risks and support both individual welfare and sustained participation in self-employment.

7. Limitations

The most important limitation is that self-employment is not strictly exogenous in the TWFE framework. While individual fixed effects control for stable personal characteristics such as risk preferences and ability, they cannot address time-varying factors which simultaneously influence income volatility and the decision to enter self-employment. A clear example is job loss: individuals who experience involuntary unemployment may be pushed into self-employment out of necessity, and this transition is also likely to be associated with increased income volatility. While a Heckman (1979) selection model was considered to address this issue, this required a valid instrumental variable affecting selection but not income volatility. No credible variable could be identified in the data, therefore this approach was not pursued (see Appendix).

A second limitation concerns income measurement. This study uses annual net labour income, which has advantages in reflecting realised living standards, but also has drawbacks. Measuring income annually means that short-term fluctuations within the year are not captured. This is particularly relevant for self-employed individuals whose income may vary substantially month-to-month due to seasonal or project-based work. Consequently, the analysis may understate true income volatility. In addition, self-reported income is subject to measurement error and potential misreporting. This issue may be more pronounced for self-employed individuals whose income can be complex to assess and report precisely. Negative income values, reflecting business losses, are included in the analysis as they represent genuine income risk. However, their inclusion makes direct comparisons with studies that focus only on positive income less straightforward.

Third, the education interaction results are limited by statistical power. In the TWFE framework, differences by education are identified from within-person changes in employment status, with education treated as a time-invariant characteristic. However, the number of switchers may still be limited, leading to imprecise estimates and wide confidence intervals. Resultantly, it is difficult to draw firm conclusions about differences across education groups.

Finally, the observation window of 2009-2019, while intentionally selected to exclude COVID-19 disruption, spans the post-financial crisis period of slow wage growth and labour market recovery. Results therefore capture relationships during a unique economic environment and may not be generalisable to periods of lower macroeconomic stress or to the post-pandemic labour market, in which self-employment patterns have shifted noticeably.

8. Conclusion

This dissertation investigates the relationship between employment status and income volatility in the UK. The central question – to what extent self-employed individuals are more exposed to income volatility than employees – is addressed through a three-stage empirical strategy moving from cross-sectional analysis, to within-person identification and finally, to an exploration of heterogeneity by occupation, education and gender.

The first stage establishes that self-employed individuals face significantly higher income volatility than employees, with a conditional gap of approximately 46 percentage points. The

second stage tightens this identification through a two-way fixed effects estimator, removing the influence of time invariant individual characteristics like risk tolerance, that simultaneously predict self-employment entry and income volatility. The within-person estimate of approximately 31 percentage points confirms the volatility premium reflects the structural consequences of employment status rather than the characteristics of those who select into it. These findings are consistent with prior UK evidence and support Hypothesis 1 of the study.

The third stage explores heterogeneity by occupation, education and gender, highlighting that the increase in income volatility associated with self-employment is not uniform across individuals. While self-employment is consistently linked to higher income volatility relative to employment, this increase is more pronounced among intermediate- and high-skilled workers than among low-skilled workers, contradicting the initial hypothesis. This suggests that higher-skilled self-employment may reflect the volatility accepted in exchange for autonomy and opportunity-driven work, consistent with a risk-return trade-off. In contrast, there is no robust evidence that the increase in income volatility associated with self-employment differs systematically by education or gender. These null findings should be interpreted with caution given the limited within-person variation available to identify heterogeneous effects in the panel.

Several limitations qualify these conclusions. The low within- R^2 across the TWFE, though expected, indicates that only a limited share of within-person variation in income volatility is explained by the model. The TWFE addresses time-invariant selection but cannot rule out time-varying confounders, most notably necessity-driven entry following job loss. The 2009-2019 observation period, while chosen deliberately to exclude COVID-19 disruption, covers a specific post-crisis labour market recovery period that may not generalise to other economic conditions.

However, the persistence of the within-person volatility gap between employment statuses suggests higher income volatility is a consistent feature associated with self-employment. This has implications for understanding the welfare consequences of the changing UK labour market. Particularly, it highlights the relevance of ongoing discussions about social protection for non-standard workers and whether existing institutional arrangements adequately account for this.

Future research could strengthen identification through a selection-correction approach with a credible exclusion restriction, addressing remaining concerns about time-varying entry into self-employment. Second, estimating separate TWFE models within education groups and using longer panels would provide a more powerful test of education heterogeneity. Finally, extending the panel beyond 2019 would allow an assessment of whether the observed within-person volatility premium associated with self-employment persists under recent shifts in labour market composition.

References

- Amit, R. and Muller, E. (1995). 'Push' and 'Pull' Entrepreneurship. *Journal of Small Business & Entrepreneurship*, 12(4), pp.64–80. Available at: <https://doi.org/10.1080/08276331.1995.10600505>.
- Audoly, R. (2024). Self-employment and labour market risks. *International Economic Review*, 66(2). Available at: <https://doi.org/10.1111/iere.12732>.
- Auguste, D., Roll, S. and Despard, M. (2022). The Precarity of Self-Employment among Low- and Moderate-Income Households. *Social Forces*, 101(3), pp.1081–1115. Available at: <https://doi.org/10.1093/sf/soab171>.
- Avram, S., Brewer, M., Fisher, P. and Fumagalli, L. (2021). Household Earnings and Income Volatility in the UK, 2009–2017. *The Journal of Economic Inequality*, 20. Available at: <https://doi.org/10.1007/s10888-021-09517-3>.
- Azariadis, C. (1975). Implicit Contracts and Underemployment Equilibria. *Journal of Political Economy*, 83(6), pp.1183–1202. Available at: <https://doi.org/10.2307/1830855>.
- Becker, G.S. (1964). *Human Capital: A Theoretical and Empirical Analysis with Special Reference to Education, First Edition*. NBER. Available at: <https://www.nber.org/books-and-chapters/human-capital-theoretical-and-empirical-analysis-special-reference-education-first-edition> (Accessed 5 January 2026).
- Benz, M. and Frey, B.S. (2008). Being Independent is a Great Thing: Subjective Evaluations of Self-Employment and Hierarchy. *Economica*, 75(298), pp.362–383. Available at: <https://doi.org/10.1111/j.1468-0335.2007.00594.x>.
- Blanchflower, D.G. and Oswald, A.J. (1998). What Makes an Entrepreneur? *Journal of Labor Economics*, 16(1), pp.26–60. Available at: <https://doi.org/10.1086/209881>.
- Blundell, R., Costa Dias, M., Joyce, R. and Xu, X. (2020). *Covid-19: the impacts of the pandemic on inequality*. Institute for Fiscal Studies. Available at: <https://ifs.org.uk/publications/covid-19-impacts-pandemic-inequality> (Accessed 2 March 2026).

Brandolini, A., Cannari, L., D'Alessio, G. and Faiella, I. (2004). *Household wealth distribution in Italy in the 1990s*. Banca D'Italia. Available at: <https://www.bancaditalia.it/pubblicazioni/temi-discussione/2004/2004-0530/index.html?com.dotmarketing.htmlpage.language=1&dotcache=refresh&dotcache=refresh> (Accessed 6 February 2026).

Cabral, A.C.G., Gemmell, N. and Alinaghi, N. (2020). Are survey-based self-employment income underreporting estimates biased? New evidence from matched register and survey data. *International Tax and Public Finance*, 28, pp.284–322. Available at: <https://doi.org/10.1007/s10797-020-09611-8>.

Cappellari, L. and Jenkins, S.P. (2014). Earnings and labour market volatility in Britain, with a transatlantic comparison. *Labour Economics*, 30, pp.201–211. Available at: <https://doi.org/10.1016/j.labeco.2014.03.012>.

Carroll, C.D. (1992). The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. *Brookings Papers on Economic Activity*, 23(2), pp.61–156. Available at: The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence | Brookings (Accessed 19 January 2026)

Carroll, C.D. (1997). Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics*, 112(1), pp.1–55. Available at: <https://doi.org/10.1162/003355397555109>.

Celik, S., Juhn, C., McCue, K. and Thompson, J. (2012). Recent Trends in Earnings Volatility: Evidence from Survey and Administrative Data. *The B.E. Journal of Economic Analysis & Policy*, 12(2). Available at: <https://doi.org/10.1515/1935-1682.3043>.

Clark, A.E. and Oswald, A.J. (1994). Unhappiness and Unemployment. *The Economic Journal*, 104(424), pp.648–659. Available at: <https://doi.org/10.2307/2234639>.

Cribb, J. and Xu, X. (2020). *Going solo: how starting solo self-employment affects incomes and well-being*. Institute for Fiscal Studies. Available at: <https://ifs.org.uk/publications/going-solo-how-starting-solo-self-employment-affects-incomes-and-well-being> (Accessed 9 January 2026).

Dahl, M., DeLeire, T., Schwabish, J. (2011). Estimates of Year-to-Year Volatility in Earnings and Households Incomes from Administrative, Survey, and Matched Data. *Journal of Human Resources*, 46(4), pp.750 – 774. Available at: <https://doi.org/10.1353/jhr.2011.0000>.

Dawson, C.J., Henley, A. and Latreille, P.L. (2009). Why Do Individuals Choose Self-Employment? *IZA Discussion Paper No.3974*. Available at: <https://ssrn.com/abstract=1336091>.

Deaton, A. (1992). *Understanding consumption*. Oxford Etc: Clarendon Press, Oxford.

Dercon, S. (2004). Growth and shocks: evidence from rural Ethiopia. *Journal of Development Economics*, 74(2), pp.309–329. Available at: <https://doi.org/10.1016/j.jdeveco.2004.01.001>.

Dynan, K.E., Elmendorf, D.W. and Sichel, D. (2007). The Evolution of Household Income Volatility. *FEDS Working Paper No. 2007 -61*. Available at: <https://dx.doi.org/10.2139/ssrn.1051141>.

Dynan, K.E., Elmendorf, D.W. and Sichel, D.E. (2012). The Evolution of Household Income Volatility. *SSRN Electronic Journal*. Available at: <https://doi.org/10.2139/ssrn.2138990>.

Evans, D.S. and Leighton, L.S. (1990). Small Business Formation by Unemployed and Employed Workers. *Small Business Economics*, 2(4), pp.319–330. Available at: <https://doi.org/10.2307/40228659>.

Freedland, M.R. and Adams-Prassl, J. (2017). Employees, Workers, and the ‘Sharing Economy’: Changing Practices and Changing Concepts in the United Kingdom. *Oxford Legal Studies Research Paper No. 19/2017*. Available at: <https://doi.org/10.2139/ssrn.2932757>.

Friedman, G.C. (2014). Workers without employers: Shadow corporations and the rise of the gig economy. *Review of Keynesian Economics*, 2(2), pp.171–188. Available at: <https://doi.org/10.4337/roke.2014.02.03>.

Friedman, M. (1957). *A Theory of the Consumption Function*. Princeton, Princeton University Press.

Giupponi, G. and Xu, X. (2020). *What does the rise of self-employment tell us about the UK labour market?* Institute for Fiscal Studies. Available at: <https://ifs.org.uk/publications/what-does-rise-self-employment-tell-us-about-uk-labour-market> (Accessed 16 January 2026).

Glavin, P., Filipovic, T. and van der Maas, M. (2019). Precarious versus Entrepreneurial Origins of the Recently Self-Employed: Work and Family Determinants of Canadians’ Self-Employment Transitions. *Sociological Forum*, 34(2), pp.386–408. Available at: <https://doi.org/10.1111/socf.12502>.

Gordon, D.F. (1974). A Neo-Classical Theory of Keynesian Unemployment. *Economic Inquiry*, 12(4), pp.431–459. Available at: <https://doi.org/10.1111/j.1465-7295.1974.tb00415.x>.

Gottschalk, P. and Moffitt, R. (2009). The Rising Instability of U.S. Earnings. *The Journal of Economic Perspectives*, 23(4), pp.3–24. Available at: <https://doi.org/10.2307/27740553>.

Gottschalk, P., Moffitt, R., Katz, L.F. and Dickens, W.T. (1994). The Growth of Earnings Instability in the U.S. Labor Market. *Brookings Papers on Economic Activity*, 1994(2), pp.217–272. Available at: <https://doi.org/10.2307/2534657>.

Hamilton, B.H. (2000). Does Entrepreneurship Pay? An Empirical Analysis of the Returns to Self-Employment. *Journal of Political Economy*, 108(3), pp.604–631. Available at: <https://doi.org/10.1086/262131>.

Heckman, J.J. (1979). Sample Selection Bias as a Specification Error. *Econometrica*, 47(1), pp.153–161. Available at: <https://doi.org/10.2307/1912352>.

Hundley, G. (2001). Why and When Are the Self-Employed More Satisfied with Their Work? *Industrial Relations: A Journal of Economy and Society*, 40(2), pp.293–316. Available at: <https://doi.org/10.1111/0019-8676.00209>.

Jacobson, L.S., LaLonde, R.J. and Sullivan, D.G. (1993). Earnings Losses of Displaced Workers. *The American Economic Review*, 83(4), pp.685–709. Available at: <https://www.jstor.org/stable/2117574?seq=1>.

Key, N., Prager, D. and Burns, C. (2017). *Farm Household Income Volatility: An Analysis Using Panel Data from a National Survey* | *Economic Research Service. USDA*. Available at: <https://www.ers.usda.gov/publications/pub-details?pubid=82563> (Accessed 24 March 2026).

Knight, F.H. (1921). *Risk, Uncertainty, and Profit*. Boston, MA: Hart, Schaffner & Marx; Houghton Mifflin Co. Available at: https://www.econlib.org/library/Knight/knRUP.html?chapter_num=1#book-reader (Accessed 18 March 2026).

Low, H., Meghir, C. and Pistaferri, L. (2010). Wage Risk and Employment Risk over the Life Cycle. *The American Economic Review*, 100(4), pp.1432–1467. Available at: <https://doi.org/10.2307/27871261>.

Lucas, R.E. (1978). On the Size Distribution of Business Firms. *The Bell Journal of Economics*, 9(2), pp.508–523. Available at: <https://doi.org/10.2307/3003596>.

McKnight, A. and Rucci, M. (2020). *The financial resilience of households: 22 country study with new estimates, breakdowns by household characteristics and a review of policy options*. London School of Economics and Political Science. Available at: https://sticerd.lse.ac.uk/CASE/_New/Publications/abstract/?index=6988 (Accessed 27 March 2026).

Meghir, C. and Pistaferri, L. (2004). Income Variance Dynamics and Heterogeneity. *Econometrica*, 72(1), pp.1–32. Available at: <https://doi.org/10.2307/3598848>.

Mincer, J.A. (1974). *Schooling, Experience, and Earnings*. NBER. Available at: <https://www.nber.org/books-and-chapters/schooling-experience-and-earnings>.

Modigliani, F. and Brumberg, R. (2005) 'Utility analysis and aggregate consumption functions: an attempt at integration', in F. Modigliani (ed.) *The collected papers of Franco Modigliani*. Cambridge, MA. The MIT Press, pp. 1-47. (Originally published 1954).

Moulton, J.G. and Scott, J.C. (2016). Opportunity or Necessity? Disaggregating Self-Employment and Entry at Older Ages. *Social Forces*, 94(4), pp.1539–1566. Available at: <https://doi.org/10.1093/sf/sow026>.

Neumann, von, Morgenstern, O. and Rubinstein, A. (1944). *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press. Available at: <https://doi.org/10.2307/j.ctt1r2gkx>.

ONS (2019). *Labour market economic commentary - Office for National Statistics*. Available at: <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/articles/labourmarketeconomiccommentary/may2019> (Accessed 11 March 2026).

ONS (2026a). *Number of People in Employment (aged 16 and over, seasonally adjusted) - Office for National Statistics*. Available at: <https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/timeseries/mgrz/lms> (Accessed 23 April 2026).

ONS (2026b). *UK Self-employment jobs SA: Total (thousands) - Office for National Statistics*. Available at:

<https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemployeetypes/timeseries/dyzn/lms> (Accessed 23 April 2026).

Parker, S.C. (2005). Explaining Regional Variations in Entrepreneurship as Multiple Occupational Equilibria. *Journal of Regional Science*, 45(4), pp.829–850. Available at: <https://doi.org/10.1111/j.0022-4146.2005.00394.x>.

Parker, S.C. (2009). *The economics of self-employment and entrepreneurship*. Cambridge: Cambridge University Press. Available at: <https://doi.org/10.1017/CBO9780511493430>.

Perry, G.E., Maloney, W.F., Arias, O.S., Fajnzylber, P., Mason, A.D., Bosch, M., and Saavedra, J. (2007). *Informality: exit and exclusion*. World Bank. Washington, DC. Available at: <https://documents.worldbank.org/en/publication/documents-reports/documentdetail/326611468163756420> (Accessed 10 April 2026).

Roy, A.D. (1951). Some Thoughts on the Distribution of Earnings. *Oxford Economic Papers*, 3(2), pp.135–146. Available at: <https://doi.org/10.2307/2662082>.

Schultz, T.W. (1961). Investment in Human Capital. *The American Economic Review*, 51(1), pp.1–17. Available at: <https://www.jstor.org/stable/1818907> (Accessed 23 February 2026).

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(7-8), pp.973–982. Available at: <https://doi.org/10.1016/j.jpubeco.2011.02.007>.

Smith, A. (1776). *The Wealth of Nations*. London: W. Strahan and T. Cadell.

Sologon, D.M. and O'Donoghue, C. (2014). Shaping Earnings Insecurity: Labor Market Policy and Institutional Factors. *Review of Income and Wealth*, 60(S1), pp.205–232. Available at: <https://doi.org/10.1111/roiw.12105>.

Storey, D.J. (1991). The Birth of New Firms: Does Unemployment Matter? A Review of the Evidence. *Small Business Economics*, 3(3), pp.167–178. Available at: <https://doi.org/10.2307/40228720>.

University of Essex (2026) *Understanding Society: Waves 1-15, 2009-2024 and Harmonised BHPS: Waves 1-18, 1991-2009*, [data collection], UK Data Service, 21st Edition, SN: 6614, Available at: <http://doi.org/10.5255/UKDA-SN-6614-22>.

Wooldridge, J.M. (2019). *Introductory Econometrics: A Modern Approach*. 7th ed. Boston, Mass.; Cengage Learning.

Ziliak, J.P., Hardy, B. and Bollinger, C. (2011). Earnings volatility in America: Evidence from matched CPS. *Labour Economics*, 18(6), pp.742–754. Available at: <https://doi.org/10.1016/j.labeco.2011.06.015>.

Appendix

Table A1: TWFE with Education subgroups

<i>DV: V_i</i>	Self-employed Coefficient	(Standard error)	N	R ²
No qualification	15.756*	(9.320)	3689	0.017
GCSE	33.612***	(4.513)	16869	0.022
A-level	31.519***	(3.711)	17998	0.024
Higher/Degree	28.644***	(2.566)	36801	0.024

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A2: Instrumental variable tests for Heckman Selection model

Variable	Coefficient	F value	Valid instrument (Yes, Yes)	
			Relevance	Exogeneity
Father's occupation	0.001**	4.42	No	No
Household size	0.001	0.66	No	No
First job employment status	0.110*	3.82	No	No
Job satisfaction	0.005***	76.91	Yes	No
Investment income	0.000	1.30	No	No
Commuting time	0.005***	160.19	Yes	No

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. 20 other variables considered were either absorbed or insignificant.