

Portfolio Choices, Firm Shocks, and Uninsurable Wage Risk

ANDREAS FAGERENG

Statistics Norway

LUIGI GUISO

EIEF

and

LUIGI PISTAFERRI

Stanford University, NBER, CEPR and SIEPR

First version received January 2016; Editorial decision February 2017; Accepted March 2017 (Eds.)

Assessing the importance of uninsurable wage risk for individual financial choices faces two challenges. First, the identification of the marginal effect requires a measure of *at least one* component of risk that cannot be diversified or avoided. Moreover, measures of uninsurable wage risk must vary over time to eliminate unobserved heterogeneity. Secondly, evaluating the economic significance of risk requires knowledge of the size of *all* the wage risk actually faced. Existing estimates are problematic because measures of wage risk fail to satisfy the “non-avoidability” requirement. This creates a downward bias, which is at the root of the small estimated effect of wage risk on portfolio choices. To tackle this problem we match panel data of workers and firms and use the variability in the profitability of the firm that is passed over to workers to obtain a measure of uninsurable risk. Using this measure to instrument total variability in individual earnings, we find that the marginal effect of uninsurable wage risk is much larger than estimates that ignore endogeneity. We bound the economic impact of risk and find that its overall effect is contained, not because its marginal effect is small but because its size is small. And the size of uninsurable wage risk is small because firms provide substantial wage insurance.

Key words: Portfolio choice, Uninsurable wage risk, Background risk, Firm shocks.

JEL Codes: G11, D1, D8

1. INTRODUCTION

How important is uninsurable wage risk for individuals’ portfolio allocations?¹ To answer this question we assemble a rich administrative household data set from Norway that allows us to overcome the identification challenges which plague most of the empirical work on the subject.

1. In the portfolio choice literature, the interest is more generally on the effect of “background” risk, *i.e.* any risk that cannot be avoided or insured. In this article, as in much of the literature, we focus on wage risk because it is the primary source of uninsurable risk faced by most individuals.

Starting with Aiyagari (1994), a large literature in macroeconomics and finance has studied how the presence of wage risk in an incomplete market setting affects the patterns of individual and aggregate savings, consumption and portfolio allocations over the life cycle, as well as the behaviour of asset prices. The theory argues that under plausible preference restrictions consumers who face uninsurable wage risk respond by accumulating precautionary savings, raising labour supply, or more generally changing the pattern of human capital accumulation (*e.g.* Levhari and Weiss, 1974). Furthermore, people reduce exposure to risks that they can avoid. In particular, they change the asset allocation of their financial portfolio by lowering the share invested in risky assets, thus tempering their overall risk exposure (Merton, 1971; Kimball, 1993; Constantinides and Duffie, 1996; Heaton and Lucas 1996, 2000b).

Motivated by these theoretical predictions and the undisputed importance for most households of labour income, one strand of research has incorporated uninsurable wage risk in calibrated models of (consumption and) portfolio allocation over the life cycle and explored its ability to reproduce patterns observed in the data (*e.g.* Heaton and Lucas, 2000b; Viceira, 2001; Cocco *et al.*, 2005; Polkovnichenko, 2007). Another strand has tried to assess the empirical relevance of wage risk in explaining portfolio heterogeneity. A fair characterization of both strands of literature is that the effect of uninsurable wage risk on portfolio allocation, though carrying the sign that theory predicts, is relatively small in size. As a consequence, this channel seems to have lost appeal as a quantitatively important determinant of household portfolio choices or as a candidate explanation for asset-pricing puzzles (such as the equity premium puzzle, see *e.g.* Cochrane, 2006).

In this article, we re-evaluate the role of uninsurable wage risk for people's willingness to bear financial risk and question the conventional wisdom of the empirical literature. We argue that this literature suffers from identification problems that also affect calibrated models of life cycle savings and portfolio allocation. Identification of the effect of uninsurable wage risk is arduous and its quantification problematic.

Identification is arduous for at least three reasons. First, in order to identify the marginal effect on portfolio choice of uninsurable wage risk one needs exogenous variation in the latter. A popular solution (Heaton and Lucas, 2000a; Angerer and Lam, 2009; Betermier *et al.*, 2012; Palia *et al.*, 2014) is to measure risk with the variance of (residual) log earnings or income typically obtained from households survey data (*e.g.* the Panel Study of Income Dynamics (PSID) in the U.S.). Another is to use second moments from subjective expectations of future incomes (Guiso *et al.*, 1996; Hochguertel, 2003) or health status (which may be particularly relevant for the elderly; Edwards, 2008). However, as a recent literature suggests, a substantial part of the residual variation in earnings is predictable and reflects individual choice rather than risk (*e.g.* Cunha *et al.*, 2005; Primiceri and van Rens, 2009; Low *et al.*, 2010; Guvenen and Smith, 2014). As for subjective expectations data, there are long-standing reservations regarding their validity and content, as well as important practical data problems: subjective expectations data are rarely available alongside longitudinal data on assets. The empirical measures described above introduce a sort of errors-in-variable problem that biases towards zero the estimated effect of risk on portfolio choice. Furthermore, as we shall discuss, the size of the downward bias can be substantial.

Secondly, notwithstanding the problem of obtaining a conceptually sound measure of uninsurable wage risk, other econometric issues may make estimates of its effect on portfolio (or other financial) choice unreliable. A key issue is that most of the evidence on the effect of wage risk comes from cross-sectional data, inducing unobserved heterogeneity bias. To give a simple example, unobserved risk aversion may determine both wage risk (through, *e.g.* occupational choice), as well as the composition of one's asset portfolio. Dealing with unobserved

heterogeneity is difficult, as one requires panel data with variation over time in wage risk, which is rare.²

A final issue is that most of the empirical literature uses survey data on assets. These are notoriously subject to measurement error and rarely sample the upper tail of the distribution (which is key, given the enormous skewness in the distribution of wealth). Moreover, both in survey and administrative data there is non-negligible censoring of stockholding because several investors choose to stay out of the stock market.

One of the contributions of this article is to develop an identification strategy that overcomes these problems. First, we isolate a component of labour income variation that truly qualifies as risk—*i.e.* one that cannot be avoided or insured. This is the component of the wage that fluctuates with idiosyncratic variation in firm performance, reflecting transmission of firm shocks onto wages. We show that this component can be used as an instrument for total residual labour income variation—which allows to deal with measurement error in wage risk. Because this component varies over time, availability of long panel data on firms and their workers makes it possible to deal with unobserved heterogeneity, thus circumventing the second obstacle to achieve reliable identification.

We use administrative data for Norway. Since Norway levies a tax on wealth, each year Norwegian taxpayers must report their assets, item by item, to the tax authority. Asset holding information is provided by third parties, implying virtually no measurement error. Moreover, the data are available for a long time span and cover the entire population, including those in the very top tail of the wealth distribution. We use these data to compute financial portfolio shares at the household level. In addition, we merge the tax records data on wealth with matched employer/employees data from the social security archives. The latter contain information on workers' employment spells and earnings in each job, as well as measures of firm performance.

Additionally, we use firm employment turnover and firm closure due to bankruptcy to construct measures of unemployment risk that complement the measure of wage risk described above.

We document a number of important findings. First, ignoring the endogeneity of wage variability but accounting for unobserved heterogeneity, we reproduce the small marginal effect of uninsurable wage risk on the portfolio allocation to risky assets that characterizes the empirical literature. However, when we instrument wage variability with the firm-variation component of wage risk, we find that the marginal effect is an order of magnitude larger. This suggests a large downward bias in prevailing estimates of the effect of uninsurable wage risk and resurrects the importance—at the margin—of wage risk for portfolio choice. In contrast, we find very small effects of unemployment risk, possibly because this type of risk is substantially insured through generous social insurance programmes in Norway.³

2. Betermier *et al.* (2012) is one exception. They deal with unobserved heterogeneity by looking at people who change industry and exploiting differences in income volatility across industries. They find that people who move from low to high volatility industries reduce exposure to stocks significantly and interpret their finding as consistent with hedging. While this marks progress, movers solve one issue but raise another: moving is endogenous and it is conceivable that the same factors that trigger mobility also affect portfolio rebalancing. While the authors show evidence that movers and stayers share similar observable characteristics, selection on unobservables (such as risk preferences) may be driving mobility. In addition, the measure of earnings volatility they use—the industry mean of the volatility of net earnings—reflects both components that qualify as risk and others that do not, as well as heterogeneity across industries. This makes it hard to estimate the economic effect of wage risk on portfolio choice.

3. Empirical estimates of the effect of uninsurable wage risk on portfolio allocations face also a problem of censoring (a large fraction of investors hold no risky assets in their portfolio). Simultaneously accounting for censoring, fixed unobserved heterogeneity, and endogeneity due to measurement error is computationally unfeasible. The very few estimators that have been proposed in the literature are based on very strong assumptions that are unlikely to hold in our specific application. Nevertheless, assuming the various biases due to unobserved heterogeneity, endogeneity of wage

Secondly, we find that the marginal effect of uninsurable wage risk varies considerably across individuals depending on their level of wealth. The portfolio response of individuals at the bottom of the wealth distribution—those with little buffers to self-insure against risk—is twice as large as that of the workers with median wealth; the effect gets smaller as wealth increases and drops to zero at the top of the wealth distribution. Uninsurable wage risk is irrelevant for those with large amounts of assets despite the fact that their compensation is more sensitive (as we document) to firm shocks. As far as we know, we are the first to document empirically the importance of wealth buffers for the effect of wage risk on portfolio choice. This helps understanding what wage risk matters for. Because low-wealth individuals are sensitive to income risk, the latter matters for explaining portfolio heterogeneity among low-wealth investors. But because the portfolio of high-wealth individuals is insensitive to income risk and because they hold the bulk of the stock market, income risk is unlikely to impact stock prices.

Finally, in assessing the economic importance of uninsurable wage risk for financial decisions one needs to separate *motive*—*i.e.* size of marginal effect—from *scope*—the size of risk itself. A full assessment of the latter would require identifying how much of the non-firm-related variation in wages is truly risk and how much is acted upon by the agent. This is hard to do in the absence of a formal model that sets out the sources of market incompleteness as well as workers' information set and corresponding economic choices. However, using the estimated parameters for the marginal effect of wage risk on the portfolio allocation, the estimated degree of firm-provided wage insurance, and a sensible estimate of the degree of predictability of workers' wage shocks obtained as a by-product of our tests, we can assess the contribution of actual wage risk to portfolio allocation. Evaluated at the sample means of these values, the effect of uninsurable wage risk is small: individuals with the average amount of wage risk have a share of risky assets in portfolio that is 1/4 of a percentage point lower than that of those facing no wage risk whatsoever. While this conclusion is similar to that of the existing literature, the economic interpretation is very different. Most papers in the literature find that the scope (the size of wage risk) is large but the motive (the causal effect of risk on portfolio choice) is small. We argue the opposite: the motive is strong—a conclusion based on our ability to isolate plausible exogenous variation in wage risk. The scope is limited primarily because firms provide workers with substantial insurance, containing considerably the size of wage risk. Because we identify separately the marginal effect of a change in background risk, the amount of insurance firms provide, and the degree of predictability of workers' wage shocks, we can run counterfactuals by altering these parameters. If firms were to provide more high-powered wage contracts (a tendency documented by Bénabou and Tirole, 2016) and start sharing shocks equally with their workers, the latter would reduce the demand for risky financial assets substantially, particularly among low-wealth workers. Equally sized changes in the degree of wage predictability would instead have a small impact on the amount of wage risk and thus on the portfolio allocation. In sum, the economic importance of uninsurable wage risk crucially hinges on the insurance role of the firm and the amount of assets available to the individual to buffer labour income shocks.

The rest of the article proceeds as follows. Section 2 reviews the empirical literature and highlights our contribution. In Section 3, we illustrate the econometric problems that arise when trying to identify the effect of uninsurable wage risk on financial decisions, and show how we tackle them. Section 4 describes the data sources. Section 5 discusses the construction of our

variance and censoring are (approximately) linear, we can gauge their sizes and obtain a back-of-the-envelope estimate of the marginal effect of uninsurable wage risk on the financial portfolio. When we do this (Section 6.1), we still find an estimate that is an order on magnitude larger than the Ordinary Least Squares (OLS) (fixed effect) estimate, implying that the key force biasing the effect of uninsurable wage risk is measurement error (*i.e.* the assumption that all residual wage variability is risk).

measures of wage risk. Section 6 turns to the estimates of the marginal effect of uninsurable wage risk on people's portfolio allocation, presents several robustness tests, and allows for wealth-driven heterogeneity in the portfolio response to wage risk. We discuss the economic effect of wage risk on the demand for risky financial assets in Section 7. Section 8 concludes.

2. LITERATURE REVIEW

Several papers provide evidence that uninsurable wage risk has a tempering effect on households portfolio allocation. In one of the first studies on the topic, Guiso *et al.* (1996) use a measure of risk obtained from the subjective distribution of future labour income in a sample of Italian workers and find that households with more spread-out beliefs of future income invest a lower share in risky assets. However, the economic effect is small: households with above average subjective earnings variance invest a 2 percentage points lower share of their wealth in stocks than households with below average uncertainty. Because they use cross-sectional data, unobserved heterogeneity cannot be controlled for.⁴ Hochguertel (2003) also relies on a self-assessed subjective measure of earnings risk available for Dutch households. The data are longitudinal, allowing him to control for unobserved heterogeneity. However, the results are similar: a negative, small effect of subjective wage income risk on the share of risky assets.

One advantage of subjective expectations is that in principle they reflect all the information available to the household; one issue, however, is that elicitation can be problematic as households may have difficulties understanding the survey question. This may result in classical measurement error as well as in households mis-reporting the probability of very low income states. Both facts are consistent with the low estimated variances of income growth compared to those obtained from panel data estimates of labour income processes. Accordingly, several papers have measured uninsurable wage risk using panel data models of workers' earnings.

Heaton and Lucas (2000a) use income data from tax records of a sample of U.S. workers to measure wage income and business income variability and correlate them with stock portfolio shares. They find a negative, but small and statistically insignificant effect of wage income variability and a negative, statistically significant but still small effect of business income variability on the demand for stocks. Unfortunately, inference is impaired both because portfolio data are imputed as well as because their measures of risk—the unconditional standard deviation of wage income and proprietary income growth—may, as we discuss in the next section, contain a large portion that reflects choice rather than risk. In addition, unobserved heterogeneity, particularly in the case of proprietary income, may be driving the results.

Angerer and Lam (2009) overcome some of these problems. They use the U.S. National Longitudinal Survey of Youth to estimate the residual variance of labour income growth, after conditioning on a number of observables. Thus, their measure of uninsurable wage risk reduces the weight of the predictable component and in addition they distinguish between transitory and permanent shocks to labour income. Perhaps because of this, compared to the previous papers they find somewhat larger effects, particularly in response to the variance of permanent shocks to labour income. Overall, a 10% increase in the standard deviation of labour income shocks lowers the portfolio stock share by 3.3 percentage points. More recently, Palia *et al.* (2014) have extended the analysis to consider several sources of risk, including labour income, returns on housing, and entrepreneurial income. They estimate that one standard deviation increase in wage risk lowers the share in stocks by 1.8 percentage points and find a larger effect on participation

4. Also using cross-sectional data, Arrondel and Calvo-Pardo (2012) find a positive correlation between subjective income risk and the portfolio risky share of French households. They argue that the result can be explained by sample selection of more risk tolerant workers into riskier occupations.

(a reduction of 5.5 percentage points). Needless to say, effects are larger when all sources of risk increase at once. Yet, because they compute uninsurable wage risk as the standard deviation of the (unconditional) growth rate of earnings, their measure likely overstates the true amount of risk people face.

Overall, this summary of the literature suggests relatively contained effects of uninsurable wage risk on the demand for risky assets. This channel has, therefore, been dismissed as an important factor in explaining portfolio allocation heterogeneity and assets prices (Cochrane, 2006; Heaton and Lucas, 2008). Yet, the likely presence of (potentially severe) measurement error in wage risk raises some doubts about this conclusion and thus on the assets prices implications. In the next section, we set up an econometric framework and argue that empirical measures of uninsurable wage risk such as those used in the literature so far are very likely to generate substantial downward biases in the marginal effect of uninsurable wage risk (and other sources of background risk). We also suggest a methodology to obtain a well-defined measure of uninsurable wage risk and a consistent estimate of its marginal causal effect.

3. ECONOMETRIC FRAMEWORK

Consider the following empirical model for the portfolio share in risky assets:

$$S_{it} = \mathbf{W}'_{it}\beta + \lambda B_{it} + r_i + \varepsilon_{it} \quad (1)$$

where S_{it} is the share of risky assets in individual i 's financial portfolio at time t , \mathbf{W}_{it} are time-varying socio-demographic characteristics related to portfolio choice (such as age and total wealth), B_{it} a measure of uninsurable background risk, r_i an unobserved individual fixed effect (which may capture heterogeneity in risk tolerance, financial and general education, or other persistent traits shifting the demand for risky assets), and ε_{it} an error term. Theory predicts $\lambda < 0$, *i.e.* people respond to unavoidable risk by reducing the amount invested in risky assets. The empirical literature has used variants of the above model, coupled with some strategy to measure risk. Success in identifying the parameter λ depends on the ability to account for the unobserved heterogeneity r_i and, as we show below, on the properties of measured uninsurable risk.

For most individuals, the key component of uninsurable risk originates from wage fluctuations. Thus, most papers assume that only source of background risk is wage risk. A general empirical strategy for measuring uninsurable wage risk consists of writing a labour earnings process such as:

$$\ln y_{ijt} = \mathbf{Z}'_{it}\gamma + v_{it} + \theta_f f_{jt} \quad (2)$$

where y_{ijt} are earnings paid to worker i by firm j at time t , \mathbf{Z}_{it} is a vector of observable wage determinants, v_{it} a component of worker's earnings volatility that is partly under the control of the agent and unrelated to the fortunes of the firm (*e.g.* unobserved changes in general human capital), and f_{jt} a firm-specific shock. The econometrician does not observe the degree of the agent's control over v_{it} . We assume that the error components f_{jt} and v_{it} are mutually uncorrelated.⁵ In keeping with the evidence below, we assume that firm shocks are passed onto wages with pass-through coefficient θ_f . We can decompose the evolution of the unobserved component of wages into two components—one that is avoidable or insurable (A_{it}), and one that is not (U_{it}). Hence:

$$\ln y_{ijt} - \mathbf{Z}'_{it}\gamma = \underbrace{(1 - \theta_v)v_{it}}_{\text{Avoidable}} + \underbrace{\theta_v v_{it} + \theta_f f_{jt}}_{\text{Unavoidable}} = A_{it} + U_{it}$$

5. Note that in most of the literature there is no information on the firm, so these two terms are conflated.

The separation of v_{it} in a component that is avoidable and one that is not (with weight θ_v) comes from recognizing that part of what the econometrician identifies as “risk” is variability in earnings that reflects, at least in part, individual choices rather than risk. For instance, time out of the labour market (inducing large swings in earnings across years) could be time invested voluntarily in human capital accumulation. Some volatility can be generated by people choosing to work longer hours, or perhaps to invest in training programmes that increase their future productivity, in response to adverse financial market shocks affecting the value of their portfolio. A recent literature suggests that a non-negligible fraction of year-to-year fluctuations in labour earnings reflect heterogeneity or choice, rather than risk (Cunha *et al.*, 2005; Primiceri and van Rens, 2009; Low *et al.*, 2010; Guvenen and Smith, 2014).⁶

In keeping with this discussion, the “true” measure of uninsurable wage risk should be:

$$\begin{aligned} B_{it} &= \text{var}(U_{it}) \\ &= \theta_v^2 \text{var}(v_{it}) + \theta_f^2 \text{var}(f_{jt}) \\ &= \rho_v V_{it} + \rho_f F_{it} \end{aligned} \quad (3)$$

where V and F are the worker-related and firm-related uninsurable risk components.

Unfortunately, this is not what is typically used in the empirical literature. First, since in survey data wages are measured with error ξ_{it} , the observed wage is:

$$\ln y_{ijt}^* = \ln y_{ijt} + \xi_{ijt}$$

Secondly, the measure of wage risk that is typically used is the overall unexplained variation in wages, *i.e.*:

$$\sigma_{it}^2 = \text{var}(\ln y_{ijt}^* - \mathbf{Z}'_{it} \boldsymbol{\gamma}) = V_{it} + \rho_f F_{it} + \sigma_{\xi}^2 = B_{it} + \varphi_{it} \quad (4)$$

where $\varphi_{it} = (1 - \rho_v) V_{it} + \sigma_{\xi}^2$. This differs from the true one because it includes the variance of the measurement error and because it assumes that the volatility of the worker component v_{it} is all unavoidable risk, while in fact a fraction $(1 - \rho_v)$ of it reflects choice-related variation.

An OLS regression of S_{it} on the measure σ_{it}^2 (omitting individual fixed effects, r_i) gives inconsistent estimates of the sensitivity of portfolio choice to wage risk.⁷ Indeed:

$$p \lim \hat{\lambda}_{OLS} = \lambda \frac{\rho_v \text{var}(V_{it}) + \rho_f^2 \text{var}(F_{it})}{\text{var}(V_{it}) + \rho_f^2 \text{var}(F_{it}) + \text{var}(\sigma_{\xi}^2)} + \frac{\text{cov}(r_i, V_{it} + \rho_f F_{it})}{\text{var}(V_{it}) + \rho_f^2 \text{var}(F_{it}) + \text{var}(\sigma_{\xi}^2)}$$

The first term is a measurement error bias: wage risk is mis-measured both because all variability in v_{it} is interpreted as risk, and because there is unaccounted noise that agents do not act upon. Furthermore, if higher risk tolerance is the only element of unobserved heterogeneity and it is associated to both less-conservative portfolios and a more volatile wage process,⁸ then

6. A predictable variation in earnings (*e.g.* a temporary reduction in hours of work due to a slowdown in demand) is not necessarily avoidable or insurable. However, the idea is that information about such event gives the ability to at least partially self-insure against it.

7. Conditional on \mathbf{W}_{it} .

8. Consider, for example, using occupation dummies to measure variation in wages, and hence risk. Empirically, the self-employed have greater year-to-year wage volatility, while public employees face lower wage and employment risk. If allocation to occupations were random, theory would predict that the high-risk types should hold more conservative portfolios than the low-risk types. But this is not what is typically found in the data. The self-employed invest more in

the second term is positive and may well exacerbate the “measurement error/conceptual risk” bias towards zero (and even produce a *positive* $\widehat{\lambda}_{OLS}$ estimate if it is large enough).⁹

In panel data, one can control for individual fixed effects. Hence, the second bias term disappears. However, the sensitivity of portfolio choice to risk remains downward biased, *i.e.*:

$$p\lim \widehat{\lambda}_{FE} = \lambda \frac{\rho_v \text{var}(\widetilde{V}_{it}) + \rho_f^2 \text{var}(\widetilde{F}_{it})}{\text{var}(\widetilde{V}_{it}) + \rho_f^2 \text{var}(\widetilde{F}_{it}) + \text{var}(\widetilde{\sigma}_{\xi}^2)} \quad (5)$$

where \widetilde{X} denotes a variable expressed in deviation from the individual mean as to remove fixed effects. The extent of the downward bias can be substantial. Even ignoring measurement error in earnings, if firms offer substantial wage insurance (*i.e.* the term ρ_f is “small”) and if a relevant share of workers-related variation in earnings is due to choice rather than to risk (*i.e.* ρ_v is small), then the FE estimate of the effect of wage risk can be much lower than the true effect.

Both conditions are likely to hold in practice. As documented by Guiso *et al.* (2005) using Italian data, firms offer partial but substantial wage insurance, implying a value of ρ_f much smaller than 1 and close to 0.01 (since their estimate of θ_f is 0.1). In Section 5, we show that this result holds also in our Norwegian data. Additionally, there is evidence that a lot of variation in individual earnings is predictable. For instance, Cunha and Heckman (2007) estimate that for U.S. skilled workers only 8% of the increase in wage variability is due to increased uncertainty and 92% to heterogeneity. Using Italian subjective earnings expectations data (which incorporate more information than that typically available to the econometrician), Kaufmann and Pistaferrri (2009) calculate that only about 1/4 of the residual earnings growth variance is risk, while the remainder is predictable variation or noise.

We take these concerns seriously and recognize that the very notion of uninsurable wage risk requires that it is exogenous and that agents have little control over it. We use firm-derived measures of wage (and employment) risk to isolate one exogenous component of the variance of individual returns to human capital and use this as an instrument for the total variance of (residual) earnings σ_{it}^2 . In the above framework, this boils down to using F_{it} as an instrument for σ_{it}^2 (while controlling for fixed effects in the risky asset share equation). Why is F_{it} a valid instrument? First, under the assumption that the firm only offers partial wage insurance to the workers (an assumption strongly supported by the evidence in Section 5), F_{it} has predictive power for σ_{it}^2 (as can be seen from equation (4) when $\rho_f \neq 0$). Secondly, once occupational sorting (or other persistent labour market traits that induce both wage volatility and shift portfolio choice) is neutralized by controlling for individual fixed effects, F_{it} is orthogonal to the residual in the portfolio allocation decision as it only reflects variability in the productivity of the firm.

stocks and have greater income volatility (see *e.g.* Georgarakos and Inderst, 2014). The “puzzle” can be explained by the fact that there is sorting into occupations based on attitudes towards risk, which confounds the impact of wage risk on portfolio choice because more risk averse individuals choose both low-risk occupations and more conservative portfolios. A similar reasoning (although producing a bias of opposite sign) applies to having traits that lead to persistently high probability of unemployment. Individuals with these traits will likely invest less in risky assets and also experience more year-to-year earnings volatility.

9. In other words, suppose that the true λ is -0.5 . If $\frac{\rho_v \text{var}(V_{it}) + \rho_f^2 \text{var}(F_{it})}{\text{var}(V_{it}) + \rho_f^2 \text{var}(F_{it}) + \text{var}(\sigma_{\xi}^2)} = 0.8$, then in the absence of unobserved heterogeneity bias, $p\lim \widehat{\lambda}_{OLS} = -0.4$. However, if $\frac{\text{cov}(r_i, V_{it} + \rho_f F_{it})}{\text{var}(V_{it}) + \rho_f^2 \text{var}(F_{it}) + \text{var}(\sigma_{\xi}^2)} = 0.3$, for example, then $p\lim \widehat{\lambda}_{OLS} = -0.1$, which exacerbates the bias even further towards zero.

It is easy to show that this strategy identifies the effect of background risk on portfolio choice as:¹⁰

$$\begin{aligned} p\lim \widehat{\lambda}_{IVFE} &= p\lim \frac{\text{cov}(\widetilde{S}_{it}, \widetilde{F}_{it})}{\text{cov}(\widetilde{\sigma}_{it}^2, \widetilde{F}_{it})} \\ &= p\lim \frac{\text{cov}(\lambda(\rho_v \widetilde{V}_{it} + \rho_f \widetilde{F}_{it}) + \widetilde{\varepsilon}_{it}, \widetilde{F}_{it})}{\text{cov}(\widetilde{V}_{it} + \rho_f \widetilde{F}_{it} + \widetilde{\sigma}_{\xi}^2, \widetilde{F}_{it})} \\ &= \lambda \end{aligned} \quad (6)$$

It is important to notice that the reduced form estimate of firm volatility onto the share of risky assets does not identify the sensitivity of the portfolio allocation to wage risk, but instead:

$$\begin{aligned} p\lim \widehat{\lambda}_{RFFE} &= p\lim \frac{\text{cov}(\widetilde{S}_{it}, \widetilde{F}_{it})}{\text{var}(\widetilde{F}_{it})} \\ &= p\lim \frac{\text{cov}(\lambda(\rho_v \widetilde{V}_{it} + \rho_f \widetilde{F}_{it}) + \widetilde{\varepsilon}_{it}, \widetilde{F}_{it})}{\text{var}(\widetilde{F}_{it})} \\ &= \lambda \rho_f \leq \lambda \end{aligned}$$

as firm shocks pass through only partially to wages. Furthermore, the difference between the true sensitivity λ and the reduced form response $\lambda \rho_f$ can be very large if firms provide substantial wage insurance, *i.e.* ρ_f is “small”. We stress this case because Hung *et al.* (2014) propose precisely this type of exercise, assigning to individual investors the stock market volatility of the firm they work for as a measure of uninsurable wage risk and estimating the portfolio response to this measure. This strategy, while similar in spirit to ours, ignores that the firm component enters with a pass-through coefficient $\rho_f < 1$. To be able to identify λ from the reduced form estimate one needs also to separately identify ρ_f . This point is missed by Hung *et al.* (2014), and their strategy would only deliver consistent estimates of λ if the worker “owned the firm”—*i.e.* in the absence of wage insurance. On the other hand, papers that use survey data sets such as the Survey of Consumer Finances (SCF) or PSID to estimate the effect of background risk on portfolio choices, cannot identify its effect as they lack matched employer-employee data to estimate F_{it} and ρ_f .

It is important to stress that we are not assuming that uninsurable risk comes only from firm-related shocks. Our exercise is simply trying to isolate a source of variation in total wage variance that is plausibly exogenous. This is all we need for the identification of the marginal effect of risk on portfolio choice. There are certainly additional sources of risk that are also exogenous, such as those associated with skill depreciation, poor health, etc., and that are independent of the firm’s fortunes. However, these are much harder to identify in our administrative data. To quantify the effect of overall risk exposure on portfolio behaviour one needs a credible estimate of the marginal effect of uninsurable risk on portfolio choice (which we have), as well as a measure of the overall level of risk. In Section 7, we propose a bounding exercise in the attempt to quantify the effect of overall risk exposure on portfolio choice.

The last econometric issue we need to address is the fact that the dependent variable is censored: a non-negligible fraction of households have no risky assets in their financial portfolio.

10. Note that a simple cross-sectional IV estimator (which ignores fixed effects) will still be inconsistent, as
$$p\lim \widehat{\lambda}_{IV} = \lambda + p\lim \frac{\text{cov}(r_i, F_{it})}{\text{cov}(\sigma_{it}^2, F_{it})}.$$

One way to handle this issue is to assume that equation (1) represents the *latent* demand for risky assets, but what is observed is a censored version of it:

$$S_{it}^c = S_{it} \times \mathbf{1}\{S_{it} \geq 0\}$$

Using a fixed effect-IV estimator when the dependent variable is censored implies that equation (6) no longer provides a consistent estimator. In principle, one could apply an estimator that deals with all three problems at once (fixed effects, endogenous regressors, and censoring of the dependent variable), such as the extension of the standard Tobit estimator considered by Honorè and Hu (2004). In practice, this estimator does not work well in our administrative large-scale data set. We will instead consider some back-of-the-envelope exercises that compare various estimators proposed in the literature to gain some knowledge about the true value of the parameter of interest λ .

In general, the data requirement for identifying the effect of uninsurable wage risk are formidable. Matched employee–employer data are needed to obtain a proper measure of (at least one component of) uninsurable wage risk; to account for individual fixed effects the data need to have a panel dimension, and the panel needs to be long enough to generate variation over time in wage risk. Finally, inference on portfolio decisions is greatly facilitated if assets and incomes are measured without error, a requirement that is rarely met in households surveys because measured incomes and financial assets are plagued with reporting error, under-reporting and non-reporting (*e.g.* Hurst *et al.*, 2015).

In the empirical analysis, we use administrative data on wages and financial assets, where measurement error is virtually absent. These data are available for over 15 years and we can identify the employer: hence, we are able to construct a measure of F_{it} that is individual—and time-varying. Because the data are a panel we can control for individual fixed effects and thus purge the estimates from unobserved heterogeneity correlated with measures of uninsurable wage risk while simultaneously driving portfolio choice (*e.g.* risk tolerance). Since we are able to simultaneously account for all the issues that plague existing empirical studies, we are giving the background risk model the best possible chance to succeed or fail and understand why it fails or succeeds.

4. DATA AND NORWEGIAN INSTITUTIONAL INSURANCE PROVISIONS

4.1. *Data*

To study whether households shelter against uninsurable wage risk by changing their risky financial portfolio, we employ high-quality data from Norway consisting of eight separate databases. All of our data are collected for administrative purposes, which substantially reduce concerns about measurement error. The data sets can be linked through unique identifiers assigned to each individual and firm in Norway (similar to SSN's and EIN's for the U.S., respectively). Here we provide a broad description of these data sets, which unless otherwise specified cover the time period 1995–2010; Appendix A.1 illustrates the features of the data in greater detail.

The *Central Population Register* contains basic end-of-year demographic information (*i.e.* gender, birth date, county of residence, and marital status) on all registered Norwegian residents. Importantly, it contains family identifiers allowing us to match spouses and cohabiting couples who have a common child. We merge this data set with information on educational attainment (from the *National Educational Database*) and information on end-of-year financial assets from tax records (*Administrative Tax and Income Register*).

To comply with the wealth tax, each year Norwegians must report to the tax authority the value of all real and financial assets holdings as of the end of the previous calendar year. Data on traded

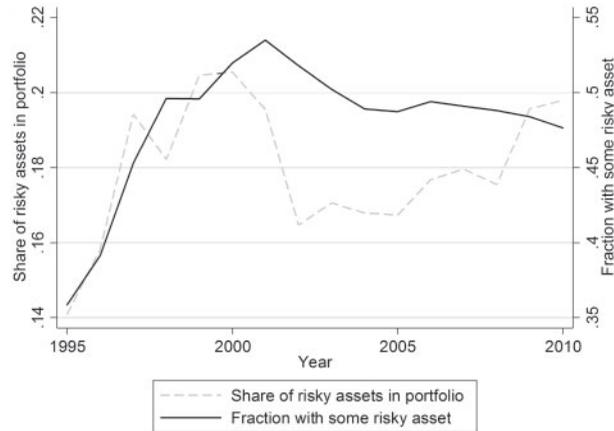


FIGURE 1

The evolution of stock market participation and the share of risky assets in household portfolios

Notes: The figure shows the average share of wealth held in risky assets (left scale) and the fraction of stockholders (right scale) among Norwegian households, by year.

financial assets, for a broad spectrum of assets categories, are reported (at their market value) directly by the financial institution that has the assets in custody (*e.g.* a mutual fund or a deposit bank). This has two main advantages: first, financial assets are measured with virtually no error; secondly, because they are reported by a third party, the scope for tax evasion is absent. For stocks of non-listed and non-traded companies, asset valuation is based on annual reports submitted to the tax authority by the companies themselves. If the tax authority finds the proposed evaluation unrealistically low, it can start a formal audit process, which limits the scope for undervaluation.

Besides the asset values data set, we have also access to the *Register of Shareholders* for the period 2004–10. This register reports, on an individual basis, the number and value of individual stockholdings, together with the ID of the firm that issues the stock. This allows us to account for direct stockholding in the company where the worker is employed, a feature that turns out to be useful when we discuss various robustness checks (Section 6.2).

Because we focus on the household as our decision unit, we aggregate assets holdings at the level of the family by summing up asset values across family members using the unique household ID described above.¹¹ We then classify financial assets holdings into “risky assets” (R)—the sum of directly held stocks in listed and non-listed companies and mutual funds with a stock component—and “risk-free assets” (RF)—the difference between total financial assets and risky assets, which includes bank deposits, government bonds, and money market funds—and define the portfolio risky assets share for each households $S_{it} = \frac{R_{it}}{R_{it} + RF_{it}}$. Because of limited stock market participation, $S_{it} = 0$ for non-participants, giving rise to censoring in our left-hand side variable.¹² In the population (before any sample selection), participation in the risky assets market increases substantially in the 1995–2010 period (Figure 1). During the same time period the average portfolio share in risky assets also increases (the dashed line in Figure 1).

11. In Norway, married couples are taxed jointly when it comes to wealth tax, but individually for income tax purposes.

12. In the original data, there are households holding extremely small amounts in stock accounts, due presumably to dormant accounts. We assume that genuine stock market participants have at least the equivalent of 30 USD worth of risky assets in their portfolio. Imposing different thresholds has no effects on the results (Table 5).

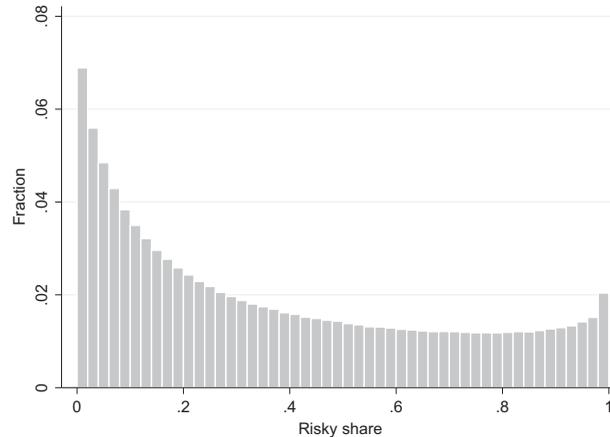


FIGURE 2

The distribution of the share of risky assets in household portfolios

Notes: The figure shows the sample distribution of the share of risky assets in the portfolio of Norwegian households. Bin size = 0.02.

Consistent with what found in the literature (Guiso and Sodini, 2013), there is substantial cross-sectional variation in the conditional risky share. As Figure 2 shows, its distribution spans the entire [0–1] range—from people holding very small amounts to people investing their entire financial portfolio in stocks. In this article, we ask how much of this heterogeneity can be explained by wage and unemployment risk, if any.

Table 1 shows summary statistics for the portfolio data and the financial wealth of our Norwegian sample. Since we select younger households with the primary earner working in the private sector (see below), their average stock market participation is higher than in the whole population (55%); conditional on participation, the average Norwegian household in our sample invests about 38% of its portfolio in risky assets.

The *Employer–Employee Register* links workers to firms; for each worker it reports all employment spells with each employer, and the compensation received. This allows us to trace the working history of each worker as she moves across firms and occupational status.

We combine the *Employer–Employee Register* with the *Central Register of Establishments and Enterprises* and the *Balance Sheet Register* with the unique firm ID present in all of these data sets. The former contains information on industry classification and institutional sector, whereas the other contains accounting data on the firm’s assets, liabilities, and income statement. Among other items, it includes data on the firm’s value added and sales that we use to construct (statistically) shocks to the firm profitability.

Lastly, on the firm side, the *Register of Bankruptcies* contains information on the date a firm enters a bankruptcy proceeding (if any) and is declared insolvent. We use this data set to identify episodes of firm closure and enrich the measure of background risk based on the variance of workers earnings with a measure of employment risk. In fact, the total variance of income comes partly from (high frequency) wage variability conditional on working, and partly from (low frequency) income variability conditional on losing the job.

Combining these three firm-level data sets with the *Employer–Employee Register* allows us to assign each worker in the sample the variability of the firm he/she works for (which depends on the pass-through coefficient estimated in Section 5), and to obtain a measure of wage risk that is theoretically more appropriate. Similarly, we can assign each worker the risk of involuntary job loss at that firm. Because our measure of risk depends on shocks to the firm that are in some degree

TABLE 1
Descriptive statistics, 1995–2010

Variable	Mean	Std. dev.	N
Age	45.518	8.545	1,972,639
Male	0.816	0.387	1,972,639
Less than high school	0.196	0.397	1,972,639
High school	0.564	0.496	1,972,639
Some college or more	0.24	0.427	1,972,639
Family size	2.881	1.405	1,972,639
Value of risky assets	479,458.1	8,854,058.0	1,972,639
Value of safe assets	359,031.3	2,253,133.8	1,972,639
Share risky assets	0.207	0.293	1,972,639
Participation share	0.552	0.497	1,972,639
Cond. share risky assets	0.375	0.304	1,089,477
Earnings	415,686.0	233,517.7	1,972,639
Earnings, family	530,971.1	306,566.9	1,972,639
Variance earnings growth	0.053	0.105	1,972,639
Variance earnings growth, family	0.077	0.127	1,972,639
Variance value added growth	0.16	0.493	284,627
Permanent shocks	0.044	0.179	205,874
Transitory shocks	0.051	0.249	243,632
Firm size	26.88	141.374	347,813
Firm bankrupt in 1 year	0.002	0.044	1,972,639
Job separation rate residual			
High education	0.003	0.192	253,017
Low education	0.001	0.194	227,402

Notes: The table shows summary statistics of households and firm characteristics. The portfolio and wealth variables, and the wage risk measures refer to our baseline sample (Appendix A.2). Asset and earnings values are NOK in 2010 prices (1 USD is approx. 5.61 NOK).

passed over to workers, we focus on a sample of individuals who are continuously employed in the private sector (30% of the workers are employed in the public sector in Norway).¹³ This excludes those who have a spell in the government sector, the retired/disabled, and those earning less than the threshold amount for unemployment benefits (the “unemployed”). We also exclude individuals who are younger than 25 years (and hence possibly still in college) and those older than 60 years (who may have intermittent participation and widespread access to early retirement). After these exclusions and a few others due to missing data at the firm level, we are left with a final sample of 4,846,766 observations. The number of observations in the various regressions we run are less than this because we use lags for constructing some of the variables and instruments. Appendix A.2 describes the sample selection in greater detail.

4.2. *Employment and wage insurance in Norway*

Portfolio (and savings) responses to wage fluctuations and risk of job loss can be affected by the extent of insurance that Norwegian workers obtain through the welfare state. For example, no matter how large the volatility of wages, portfolio choice would be independent of it if wage risk were fully insured.

Here we provide a broad description of social insurance programmes in Norway, which are indeed relatively generous by international standards. First, workers receive unemployment insurance (UI). For permanent layoffs, UI lasts for 52–104 weeks and replaces, on average, 62%

13. If there are multiple earners in the household (and both work in the private sector), we measure wage risk with the one faced by the primary earner.

of the gross income in the last occupation (up to a cap). For temporary layoffs, UI is limited to 26 weeks within a 1.5-year period since layoff. Norway offers also disability insurance, which is obtained when the assessed loss in earnings capacity is at least 50%. Unlike the U.S., eligibility is means tested (based on income and assets). Finally, individuals may have access to sickness and maternity benefits and active labour market programmes to revamp their skills in case of displacement.

While Norwegian workers are better shielded than, say, U.S. workers against extreme low realizations of their human capital (*i.e.* their consumption floor is higher), they do face substantial uninsured risk. First, while UI is generous (at least relative to U.S. standards), unemployment risk is not fully insured: UI benefits are time limited, replace a fraction of lost wages, and remaining unemployed is economically costly due to scarring effects (Nilsen and Reiso, 2014). Indeed, despite the institutional differences, in the 2001–13 period average duration of unemployment in Norway was only 15% longer than in the U.S. for people aged 25–54 years.¹⁴ Secondly, there is little government protection against the risk of wage fluctuations conditional on employment—especially those induced by firm-related shocks. There is indeed no insurance against wage cuts or *not* receiving bonuses, but there is against being laid off. Finally, while severe wage fluctuations induced by, say, health limitations are insured through the disability insurance system, the means-tested aspect of the programme reduces the scope of insurance, in particular due to the relative low risk of a disability and the fungibility of savings (*e.g.* due to retirement or bequest motives).

5. MEASURING RISK

In this and following sections, we discuss our empirical findings. We start by motivating economically our instruments. Next, we estimate the marginal effect of uninsurable wage risk on portfolio allocation. Finally, we assess the robustness of our findings.

To construct a measure of wage risk that can be arguably considered as unavoidable or uninsurable, we focus on shocks to firm profitability, which may induce variation in workers' pay (conditional on retaining the job) or even involuntary job loss in more extreme cases. This strategy requires that: (1) we measure firm-related shocks; and (2) we identify how much of these shocks are passed onto the worker's wages.

In principle, our instrument would be economically irrelevant if labour markets were frictionless and workers could move rapidly and without cost between firms. A frictionless labour market would, effectively, provide them with full insurance against firm idiosyncratic shocks. The fact that firm shocks are passed onto wages (as we document below) is of course *prima facie* evidence against this possibility.

The idea that firm-specific shocks are passed onto workers' earnings requires that wages are at least partly determined at the firm level. This in turn depends on the wage-setting process. In Norway, like in other Nordic countries, union density and coverage are high. However, in the private sector the coverage of collective bargaining agreements is actually “only” 55%, leaving ample room for many workers to have wages set outside the conventional framework. Even for workers whose wages are negotiated centrally, there is still ample room for local negotiation (or wage drift). Moreover, for white collars, collective bargaining only determines the procedures for setting wages, while the actual level of wages is negotiated on an individual basis. Finally, as reported by Loken and Stokke (2009), the share of private sector employees with a component of pay that is variable (and most likely related to the firm performance) has increased considerably from 10% in 1990 to 40% in 2005.

14. See OECD statistics at <http://stats.oecd.org>. “Labour”—“Labour Force Statistics”—“Unemployment by duration”.

5.1. Wage risk: firm shocks and pass-through

Following Guiso *et al.* (2005), we measure firm j 's performance with its value added, VA_{jt} ,¹⁵ and assume its log evolves according to the process:

$$\begin{aligned}\ln VA_{jt} &= \mathbf{X}'_{jt}\varphi + Q_{jt} + f_{jt}^T \\ Q_{jt} &= Q_{jt-1} + f_{jt}^P\end{aligned}$$

where \mathbf{X}_{jt} is a vector of observables that captures the predictable component of firm's performance. The shock component is the residual $Q_{jt} + f_{jt}^T$, the sum of a random walk component Q_{jt} with permanent shock f_{jt}^P and a transitory shock component f_{jt}^T .

Next, we model the earnings y_{ijt} (in logs) of worker i in firm j , in a similar vein, as a linear function of a predictable component that depends on a vector of workers observed characteristics, \mathbf{Z}_{ijt} , an individual random walk and transitory component, and a component that depends on the firm shocks with transmission coefficients θ^T and θ^P , respectively, for transitory and permanent firm value added stochastic component.¹⁶ Hence we generalize (2) and write:

$$\begin{aligned}\ln y_{ijt} &= \mathbf{Z}'_{ijt}\gamma + v_{it} + \theta^P Q_{jt} + \theta^T f_{jt}^T \\ v_{it} &= P_{it} + \eta_{it} \\ P_{it} &= P_{it-1} + \chi_{it}\end{aligned}$$

For firm-related uninsurable wage risk to matter, θ^T and θ^P must be positive and significant. That is, firms must pass over to the workers some of the shocks to their performance and not offer them full wage insurance. Using Italian data, Guiso *et al.* (2005) show that firms offer partial wage insurance with respect to permanent and transitory shocks—that is the estimated values of θ^T and θ^P are positive but smaller than one—and that the pass-through is larger for permanent shocks. Replicating their methodology, their result has been shown to hold also in other countries, such as Portugal (Cardoso and Portela, 2009), Germany (Guertzgen, 2014), Hungary (Katay, 2009), Sweden (Friedrich *et al.*, 2015), Belgium (Fuss and Wint, 2008), France (Biscourp *et al.*, 2005), and the U.S. (Lagakos and Ordonez, 2011; Juhn *et al.*, 2017) with remarkably similar patterns.

To establish the degree of pass-through of firm shocks to wages in Norway, we use Guiso *et al.*'s (2005) methodology. Define the unexplained growth of firm value added, g_{jt} , and of workers' earnings, ω_{ijt} as:

$$\begin{aligned}g_{jt} &= \Delta(\ln VA_{jt} - \mathbf{X}'_{jt}\varphi) \\ \omega_{ijt} &= \Delta(\ln y_{ijt} - \mathbf{Z}'_{ijt}\gamma)\end{aligned}$$

15. Firm value added is defined as revenues minus operating costs other than labour and capital costs (*i.e.* materials and services, such as rents, advertisement, and R&D). This is a standard measure of firm productivity or performance that captures the total economic value created (added) by the joint use of capital and labour employed within the firm.

16. These processes fit the data quite well. The first-order autocovariances in the residual of the wage equation and in the firms value-added equation are negative, economically large and highly statistically significant. The higher-order autocovariances decay very rapidly (the second order autocovariance is 10 times smaller than the first order one in both processes). Not surprising given the very large number of observations, they retain statistical significance. Economically, however, autocovariances past the second lag are minuscule.

TABLE 2
Pass-through of firms' shocks to workers' wages

	Permanent value added shocks	Transitory value added shocks
	(1)	(2)
Pass-through coef.	0.0705*** (0.0056)	0.0175*** (0.0053)
Constant	-0.0021*** (0.0002)	-0.0023*** (0.0002)
Hansen <i>J</i> -test <i>p</i> -value	0.54	0.00
<i>F</i> -stat first stage	134.21	688.46
Observations	2,358,889	2,370,420

Notes: The table reports estimates of the pass-through coefficient of permanent (column (1)) and transitory shocks (column (2)) to the firm's performance onto its workers' wages, using the identification strategy of Guiso *et al.* (2005). Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

Guiso *et al.* (2005) show that the pass-through coefficients θ^T and θ^P can be identified by simple IV regressions:

$$\theta^T = \frac{\text{cov}(\omega_{ijt}, g_{jt+1})}{\text{cov}(g_{jt}, g_{jt+1})}$$

$$\theta^P = \frac{\text{cov}(\omega_{ijt}, g_{jt-1} + g_{jt} + g_{jt+1})}{\text{cov}(g_{jt}, g_{jt-1} + g_{jt} + g_{jt+1})}$$

Accordingly, we preliminarily run regressions for firm value added and workers' wages. In the first, we control for year dummies, area dummies, sector dummies, log firm size, and in the second for year dummies, a quadratic in age, dummies for the quantity and type of schooling, firm size, dummies for whether the individual experienced periods out of work due to sickness, maternity leave, or unemployment, family size, area dummies, dummies for immigration status, and for family type. We then retrieve the residuals from these regressions (the empirical analogs of g_{jt} and ω_{ijt} above), and estimate θ^T and θ^P . Results for the pass-through estimates are shown in Table 2.¹⁷

Both parameters θ^T and θ^P are positive and estimated with great precision, implying that both permanent and transitory shocks to the firm value added are passed onto wages. As in Guiso *et al.* (2005), the wage response to permanent shocks to the firm performance (0.071) is significantly larger than the response to transitory shocks (0.018), which accords with intuition. The value of the *F*-test suggests that the instruments used to identify the two parameters are quite powerful while the Hansen *J*-test of the overidentifying restrictions reveals some misspecification for θ^T , possibly arising from the fact that the i.i.d. assumption is a bit restrictive. Given that transitory shocks play a small role, this is not worrying.

17. If workers respond to a negative firm permanent shock by working permanently more, then our estimate of θ^P is downward biased (there is actually less insurance than we estimate). How large is the bias? This type of reaction requires that workers have high Marshallian elasticities, but for men these are typically very small, so we conjecture that the bias is not very large. We also suspect that not many workers can work longer hours when the firm is distressed (conditional on staying).

To have a reasonably long series of wage volatility measures, our strategy is to compute the overall variance of unexplained workers earnings growth over T periods using rolling averages¹⁸:

$$\sigma_{it}^2 = \frac{\sum_{s=0}^{T-1} \omega_{ijt-s}^2}{T} \quad (7)$$

We use this measure as explanatory variable when estimating the risky portfolio share but instrument it with the variances of the unexplained firm value added growth—both permanent and transitory—computed over the same T periods:

$$F_{jt}^P = \frac{\sum_{s=0}^{T-1} g_{jt-s}(g_{jt-s-1} + g_{jt-s} + g_{jt-s+1})}{T} \quad (8)$$

$$F_{jt}^T = \frac{\sum_{s=0}^{T-1} g_{jt-s}g_{jt-s+1}}{T} \quad (9)$$

Notice that since the computation of these variances requires using lagged values of growth rates, it can only be implemented if the panel has a long time dimensions, which is the case in our data. We set $T=5$ in what follows.¹⁹

5.2. Unemployment risk

Our second measure of background labour income risk is unemployment risk. This risk should also in principle reflect idiosyncratic shocks to the (worker's) firm so that it can vary across workers and over time.²⁰

We construct two indicators. The first is job loss risk induced by firm closure. We use the Registry of Firm Bankruptcies, which records the date in which the firm is declared insolvent. We construct an indicator of firm closure risk if the worker is currently working in a firm that will be declared bankrupt in t years, and experiment by changing the lead value t . The usefulness of this measure is that job losses arising from firm closures are involuntary (and hence unavoidable).

We also experiment with an alternative proxy of job loss risk based on the idea that workers employed in more volatile firms (in terms of workforce turnover) face a greater likelihood of separation. To allow for the fact that, even within volatile firms, workers of different skills may be exposed to different risk of job loss, we construct this indicator separately for low- and high-education workers. To build this measure, we start by constructing an indicator of firm/education-specific job separation:

$$JS_{jqt} = \frac{F_{jqt}}{S_{jqt-1}} \quad (10)$$

where F_{jqt} is the separation flow of workers of education q from firm j at time t , and S_{jqt-1} the corresponding stock at time $t-1$. There are two remarks about this measure: (1) unlike what

18. This measure can be interpreted as an approximation to a conditional variance prediction one would get from estimating an Autoregressive Conditional Heteroskedasticity (ARCH)(T) process (with restrictions on the ARCH coefficients). Meghir and Pistaferri (2004) find evidence that wage risk can be modelled as an ARCH process. The assumption is that individuals form forecasts about risk using recent realizations of the process (squared) innovations.

19. The results are qualitatively similar if we use $T=3$ or $T=4$.

20. Unemployment risk arising from macroeconomic fluctuations in economic activity constitutes uninsurable risk but, being common to all workers, is of little help in identifying the effect of labour income risk on financial decisions.

done in most of the literature, availability of worker-level data allows us to express the separation flow F_{jqt} as gross rather than net employment changes; (2) not all separations reflect involuntary layoffs (some are voluntary quits or reflect events such as retirement, migration, intra-group mobility, etc.)—and hence equation (10) is only a (noisy) proxy of unemployment risk. To isolate “excess job separations”, we regress equation (10) on firm fixed effects and dummies for year and size, and take the residuals from this regression as our proxy of idiosyncratic job loss risk. This procedure captures how job separation at the current firm differs from that of a similar firm (in terms of size and business cycle) and relative to its long-term behaviour (since firms do not change industry denomination, the fixed effect component also accounts for industry-specific effects).

The bottom part of Table 1 reports summary statistics for our measures of labour market risk along with the estimated variances of the firms shocks. We find that the average variance of earnings growth in our sample is 0.05, with a standard deviation of 0.11; both figures are small compared to those estimated from survey data (*e.g.* Gourinchas and Parker, 2002; Cocco *et al.*, 2005) partly reflecting absence of measurement error in our measure of earnings. In contrast, the variance of firm value added growth is much larger (0.16), with an extremely large standard deviation of 0.49. Finally, the last three rows report measures of the risk of firm closure and of job loss based on firm turnover. The first risk appears contained—the average probability of firm closure is 0.2%. However, the consequences of involuntary job loss associated with firm bankruptcy may be quite disastrous, at least for some workers, due to scarring effects.²¹ Allowing for job loss risk we can study the role of idiosyncratic tail background risk in households financial decisions whose importance for assets pricing has been recently stressed by Schmidt (2016).²² The last rows in the table report descriptive statistics on our second proxy of job loss risk (firm-specific excess job separation rates). These rates are centred around zero (by design) and their standard deviation is around 19%. There are no significant differences between education levels.

6. THE EFFECT OF BACKGROUND RISK ON THE RISKY PORTFOLIO SHARE

In this section, we test whether and by how much investors mitigate uninsurable risk in their human capital by reducing financial risk (through rebalancing of their financial portfolio away from stocks). We start with regressions of the portfolio share of risky financial assets against a set of socio-demographic characteristics of the household, our measure of wage risk—the variance of unexplained wage growth—and households fixed effects to capture general heterogeneity in preferences for risk that can be correlated with uninsurable wage and unemployment risk. Of course, these fixed effects may also capture other sources of unobservable heterogeneity that may impact households portfolio allocation—such as differences in the precision of information about stock returns (Peress, 2004) or financial sophistication (Calvet *et al.*, 2009).

Our empirical specification includes a rich set of controls: a quadratic in age to model life cycle portfolio effects, year dummies which may capture passive variation in the asset share in response to common changes in stock prices, and dummies for family type and area of residence. To capture well-documented differences in assets allocation due partly to fixed participation costs

21. Nilsen and Reiso (2014) study the long-term unemployment consequences of displacement in Norway. They find that 5 years after job displacement, the likelihood of being unemployed is still 17.2% among the “treated” group and only 7.8% among the “control” group. The negative effect decreases over time, but there is some unemployment “scarring” effect remaining even 10 years after the initial shock.

22. Calibrated life-cycle portfolio models find small effects of uninsurable wage risk on the portfolio share in stocks but larger effects, particularly at young age, for the idiosyncratic risk of a job loss associated with a large wage cut (Viceira, 2001; Cocco *et al.*, 2005). However, this latter effect is obtained assuming no UI.

TABLE 3
The effect of labour income risk on the financial portfolio share

	(1)	(2)	(3)
	Fixed effect	Reduced form fixed effect	Fixed effect IV (<i>Baseline</i>)
σ_{it}^2	-0.0202*** (0.0029)		-0.4986*** (0.1827)
F_{it}^P		-0.0033*** (0.0012)	
F_{it}^T		-0.0028*** (0.0007)	
Firm bankrupt in 1 year		-0.0112** (0.0050)	-0.0201*** (0.0066)
Firm bankrupt in 3 years		0.0008 (0.0027)	-0.0040 (0.0034)
Firm bankrupt in 5 years		0.0006 (0.0028)	-0.0020 (0.0035)
Lagged log wealth	0.0153*** (0.0002)	0.0104*** (0.0002)	0.0112*** (0.0003)
Home ownership	0.0176*** (0.0009)	0.0135*** (0.0009)	0.0146*** (0.0012)
Age	0.0224*** (0.0006)	0.0195*** (0.0005)	0.0222*** (0.0017)
Age sq.	-0.0002*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)
Family size	-0.0005 (0.0006)	0.0005 (0.0006)	-0.0013* (0.0008)
Hansen <i>J</i> -test <i>p</i> -value			0.13
<i>F</i> -stat first stage			56.85
Observations	1,972,639	1,655,104	1,184,800

Notes: The table reports estimates of the marginal effect of wage and unemployment risk on the risky financial portfolio share. Column (1) shows fixed effect regressions; column (2) reports reduced form regressions of the share on the two instruments—the variance of transitory and permanent shocks to firm's value added. Column (3) shows IV estimates. Regressions also control for family type, area, and year dummies. Hansen *J*-test for instrument validity and *F*-stat for the power of the instruments are shown at the bottom of the table. Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

in the stock market and financial sophistication (Campbell, 2006), we control for lagged wealth. To account for interactions between levels of stockholding and housing (Cocco, 2004), we also control for homeownership status. For the time being, we neglect the censoring issue, which we deal with in the next section. Results of these estimates are shown in column (1) of Table 3.

The estimated coefficient on σ_{it}^2 is consistent with the idea that workers who face unavoidable wage risk tend to take less financial risk. The effect of wage risk is negative and very precisely estimated. However, its size is small: one standard deviation increase in the (residual) variance of log earnings would reduce the risky assets share by 0.12 percentage points. Because the average risky assets share over the sample period is 21%, this amounts to 0.6% of the average sample share, too small an effect to matter. Hence, these estimates replicate the small economic effect of uninsurable wage risk that has been found in the literature.

The second column shows results of the reduced form regression for the share where the reduced form instruments are the firm permanent and transitory variance of firms value added, and find again negative coefficients and much smaller responses. As argued in Section 2, this is consistent with the estimated effect of the variance of firm value added being the product of the true response of the share to wage risk and the effect of firms variability on the latter (typically

considerably smaller than 1, as shown in Table 2). Because of this, a regression of the share on the variance of firm performance does not identify the marginal effect of uninsurable wage risk.

Estimates change considerably when we instrument the wage growth variance with the permanent and transitory variance of firm performance (column 3). The coefficient on the worker's wage growth variance is negative and highly statistically significant and its size (in absolute terms) increases by a factor of 25—from -0.02 to -0.5 , resulting in a very high sensitivity of portfolio decisions to uninsurable wage risk. Of course, the economic importance of wage risk depends both on its marginal effect as well as on the size of uninsurable wage risk. In Section 7, we discuss how we can use our empirical strategy to obtain a plausible measurement of the latter and hence assess the economic contribution of wage risk.

In the regressions reported in columns (2) and (3), we also include our preferred measure of unemployment risk (firm closure). We find that the risk of firm closure discourages investment in risky assets (and the effect increases as we get closer to the firm closure event, which conforms with intuition). But the marginal effect is small. A 10-fold increase in the risk of firm closure from the sample average would reduce the share invested in risky assets by 0.2 percentage points, about 1% of the sample mean share. The fact that workers reduce stock exposure in anticipation of firm closure is *prima facie* evidence that they perceive this risk or the general distress of the firm. One may wonder whether the response we document is small because workers avoid the risk they face by abandoning in advance the “sinking ship” and smoothly relocating to another firm. To assess this possibility we estimate a probit model for the event of job mobility as a function of current and future firm shocks and worker's socio-demographic characteristics (the results are reported in Appendix A.3, Table A1). We find that future shocks to the firm growth and indicators for whether the firm goes bankrupt within 1–2 years have no statistically significant effect on mobility, implying that there is no support for the idea that workers “leave the ship before it sinks”. The fact that workers adjust their investments in stocks in response to plant closure but do not relocate is consistent with the idea that mobility is costly to implement and that protection against job loss through labour market relocation is hard to come by due to frictions.

6.1. *Dealing with censoring*

The estimates in Table 3 address two of the issues that identification of the effect of uninsurable wage risk poses—unobserved heterogeneity and the endogeneity problems that characterize the measures of risk used in the literature. The third problem, neglected so far, is that about half of our sample is censored from below at 0, *i.e.* there are on average about 45% stock market non-participants.

A formal treatment of censoring (*e.g.* through a Tobit approach) is unfeasible because we have to deal simultaneously with three issues: endogeneity of the background risk measure, unobserved heterogeneity in risk preferences that we capture with fixed effects, and censoring. Honorè and Hu (2004) propose an estimator that deals with these three issues at once, but their estimator is based on strong assumptions. For example, it requires that the endogenous variable is bounded from above and below (which in our case, where the endogenous variable is a variance, clearly is not).

Nevertheless, we can get a sense of the relative importance of the three issues for the estimates of the effect of background risk on the portfolio allocation by comparing five different models: (1) Linear regression with households FE; (2) IV linear regression with households fixed effects (IVFE) (both of which we have already discussed in Table 3); (3) IV linear regression in which we replace the fixed effects with a rich control function strategy that includes observable fixed heterogeneity (IVC); (4) IV Tobit regression with the same control function (IVTC); and (5) a

TABLE 4
Assessing the relevance of unobserved heterogeneity and censoring

	(1)	(2)	(3)	(4)
	Fixed effect IV (<i>Baseline</i>)	IV with control function	Tobit IV with control function	Tobit w/ double control function
σ_{it}^2	-0.4986*** (0.1827)	-0.4144*** (0.1152)	-0.3199*** (0.1806)	-0.373* (0.1990)
Firm bankrupt in 1 year	-0.0201*** (0.0066)	-0.0032 (0.0069)	-0.0157 (0.0125)	-0.0211 (0.0134)
Firm bankrupt in 3 years	-0.0040 (0.0034)	0.0037 (0.0047)	0.0025 (0.0085)	0.0030 (0.0083)
Firm bankrupt in 5 years	-0.0020 (0.0035)	0.0055 (0.0048)	0.0057 (0.0089)	0.0074 (0.0085)
Lagged log wealth	0.0112*** (0.0003)	0.0535*** (0.0003)	0.1187*** (0.0005)	0.0614*** (0.0007)
Home ownership	0.0146*** (0.0012)	0.0248*** (0.0011)	0.0553*** (0.0021)	0.0190*** (0.0026)
Age	0.0222*** (0.0017)	0.0141*** (0.0011)	0.0235*** (0.0017)	0.0153*** (0.0017)
Age sq.	-0.0001*** (0.0000)	-0.0002*** (0.0000)	-0.0003*** (0.0000)	-0.0004*** (0.0000)
Family size	-0.0013* (0.0008)	0.0076*** (0.0007)	0.0143*** (0.0011)	0.0056*** (0.0016)
Male		0.0294*** (0.0016)	0.0298*** (0.0027)	0.0292*** (0.0027)
Education length indicators ^a		799.88 [0.0000]	945.86 [0.0000]	576.46 [0.0000]
Education type indicators ^a		628.90 [0.0000]	426.90 [0.0000]	467.18 [0.0000]
Hansen <i>J</i> -test <i>p</i> -value	0.13	0.05	0.08	
<i>F</i> -stat first stage	56.85	347.91	347.91	
Observations	1,184,800	1,230,063	1,230,063	1,230,063

Notes: The table reports estimates of the marginal effect of wage and unemployment risk on the risky financial portfolio share. Column (1) reproduces the IV benchmark regression of Table 3, column (3); column (2) shows IV estimates but replaces the fixed effect with a control function; column (3) shows Tobit IV estimates using the same control function; column (4) shows Tobit estimates with double control functions. All regressions include controls for family type, year, and area dummies. Regressions in columns (2)–(3) add controls for gender, as well as education length and type indicators. Further, column (4) includes also means of the control variables at the individual level, and the residual from an OLS-regression of σ_{it}^2 on the reported control variables, F_{it}^p and F_{it}^T , as well as the individual means of the latter. Hansen *J*-test for instrument validity and *F*-stat for the power of the instruments are shown at the bottom of the table. In the third column the reported *p*-value for the test for instrument validity comes from a two-step procedure for computational reasons. Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

^aReported are the chi square-statistics for a test of joint significance of, respectively, 9 education length and 9 education type indicators, *p*-values in square brackets.

“double control function” estimator (2IVTC), in which one assumes a linear relationship between the fixed effect and the endogenous covariates, as in Chamberlain (1984).

If the three issues (endogeneity, fixed effects, and censoring) are all important (and if the relationship between the fixed effect and the endogenous covariates takes a more general form), none of these models delivers consistent estimates. However, the bias of each of these models is different and can potentially be compared—as we do below—to gauge their relative importance and thus enable us to say something about the true value of λ . The Online Appendix provides a discussion of the different biases.

We have already shown estimates for models (1) and (2) in Table 3 and reproduce the results of equation (2) in the first column of Table 4. In the second column, we drop the fixed effects and replace them with a rich control function that now includes nine dummies for the length

of education, nine dummies for the type of education, plus a dummy for the gender of the household head (which are very key determinants of risk tolerance or financial sophistication, see Guiso and Sodini, 2013). The estimate of λ drops (in absolute value) from -0.5 to -0.41 (which is consistent with the idea that omission of fixed effects generates an upward bias, *e.g.* because more risk tolerant investors select jobs with higher firm volatility). Though relatively large, this is not a dramatic drop from a qualitative point of view, an indication that the upward bias from omitting fixed effects is likely contained (at least conditioning on the rich control function). Column (3) shows estimates of a formal Tobit IV model with the same control function as in column (2), which should eliminate the bias from neglecting censoring. The estimate of λ is smaller but in the same ballpark, -0.32 . The difference between IVTC and IVC can be interpreted as the bias induced by censoring.²³

In the final column (4) we implement a “double control function” estimator.²⁴ In a first step, we follow Blundell and Smith (1986), run a regression of our endogenous variable σ_{it}^2 on the (included and excluded) instruments and their means (to account for individual fixed effects in the wage variances, as suggested by Chamberlain, 1984), and save the residuals, \widehat{e}_{it} .²⁵ In a second step, we run a Tobit regression on σ_{it}^2 , the residual \widehat{e}_{it} , the exogenous covariates W_{it} , and their means (to account for individual fixed effects in the risky share equation). While the estimate is noisier due to the addition of many covariates, the size of the coefficient estimate is very similar, confirming the general pattern of results.

The fact that the IVFE, IVC, IVTC, and 2IVTC estimates are of the same order of magnitude while the FE estimate is an order of magnitude smaller (in absolute value), suggests that the biases from ignoring censoring or unobserved heterogeneity are sizeable but comparatively much smaller than the endogeneity bias. What is key is accounting for the latter.

6.2. Robustness

In this section, we discuss various robustness analyses and extensions.

Variable definition: Our main indicator of job loss risk (firm closure) may capture only tail events. We hence consider the alternative measure described in Section 5.2 (firm/education job separation risk, see equation (10)). The results are reported in Table 5, column (2). We find that this alternative measure also discourages investment in risky assets, but the marginal effect remains small. A standard deviation increase in unemployment risk would reduce the share invested in risky assets by only 0.2%, again too small an effect to matter.

Our wage variance measure (7) weights all past realizations equally. However, the information contained in more recent events may be more relevant for forming forecasts about wage risk. We hence experiment by constructing rolling averages that weight recent observations more than distant observations (and do the same also for the firm variances (8) and (9)). In particular, we use declining weights, such that the realization at time t is weighted twice as much as that of time $t-1$. Results with this alternative weighting scheme are reported in Table 5, column (3). The marginal effect is somewhat smaller (-0.42), but the broad qualitative picture is unchanged.

23. Since the Tobit model is nonlinear while all the other models are linear, the bias induced by omitting fixed effects is different for the IVTC and IVC estimators. Hence, the difference between the two estimators reflects both censoring and the different incidence of fixed effects bias. We assume the latter difference is small.

24. We thank Francis Vella for suggesting this approach.

25. In other words, we assume that $\sigma_{it}^2 = z'_{it}\theta + m_i + \varepsilon_{it}$. Chamberlain (1984) suggests to model the fixed effect m_i as $m_i = z'_{i0}a_0 + \dots + z'_{iT}a_T + l_i$. To reduce the computational burden, we assume instead $m_i = \bar{z}'_i a + l_i$.

TABLE 5

Robustness: including individual-specific unemployment risk, quadratic weights, and varying threshold for risky asset market participation

	Fixed effect IV (<i>Baseline</i>) with				
	Fixed effect IV (<i>Baseline</i>)	unemployment risk	quadratic weights	risky assets > 50 USD	risky assets > 100 USD
	(1)	(2)	(3)	(4)	(5)
σ_{it}^2	-0.4986*** (0.1827)	-0.5052*** (0.1830)	-0.4159*** (0.1692)	-0.4922*** (0.1826)	-0.5002*** (0.1830)
Firm bankrupt in 1 year	-0.0201*** (0.0066)	-0.0178*** (0.0067)	-0.0152** (0.0070)	-0.0200*** (0.0066)	-0.0196*** (0.0066)
Firm bankrupt in 3 years	-0.0040 (0.0034)	-0.0042 (0.0034)	-0.0021 (0.0034)	-0.0039 (0.0034)	-0.0035 (0.0034)
Firm bankrupt in 5 years	-0.0020 (0.0035)	-0.0021 (0.0035)	-0.0022 (0.0036)	-0.0019 (0.0035)	-0.0018 (0.0035)
Lagged log wealth	0.0112*** (0.0003)	0.0112*** (0.0003)	0.0113*** (0.0003)	0.0112*** (0.0003)	0.0113*** (0.0003)
Home ownership	0.0146*** (0.0012)	0.0146*** (0.0012)	0.0145*** (0.0013)	0.0145*** (0.0012)	0.0146*** (0.0012)
Age	0.0222*** (0.0017)	0.0222*** (0.0017)	0.0241*** (0.0013)	0.0223*** (0.0017)	0.0222*** (0.0017)
Age sq.	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)
Family size	-0.0013* (0.0008)	-0.0013* (0.0008)	-0.0130*** (0.0008)	-0.0013* (0.0008)	-0.0013 (0.0008)
Unemployment risk		-0.0090*** (0.0012)			
Hansen <i>J</i> -test <i>p</i> -value	0.13	0.1234	0.1337	0.1286	0.1111
<i>F</i> -stat first stage	56.85	56.79	45.98	56.84	56.84
Observations	1,184,800	1,184,800	1,184,800	1,184,800	1,184,800

Notes: The table reports estimates of the marginal effect of wage and unemployment risk on the risky financial portfolio share. Column (1) reproduces the IV benchmark regression of Table 3; column (2) shows the results including a proxy for individual-specific unemployment risk (the likelihood of separation from a job) estimated separately for high- (at least high school) and low-education (less than high school) workers within the firm; column (3) shows the results when weighting the terms in the lag structure of the equations for σ_{it}^2 , F_{it}^p , and F_{it}^T (Section 5.1) using a quadratic mean, where $1.6/(0.1+0.2+0.4+0.8+1.6)$ is the weight on the most recent observation. Columns (4) and (5) display the estimates when varying the threshold for risky asset market participation to 50 and 100 USD, respectively. Regressions also control for family type, area, and year dummies. Hansen *J*-test for instrument validity and *F*-stat for the power of the instruments are shown at the bottom of the table. Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

Finally, we have assumed that stock market participants have at least the equivalent of 30 USD worth of risky assets in their portfolio. In columns (4) and (5), we use higher thresholds (50 and 100 USD). The results are unchanged.

Instrument validity: Our instruments for the workers' unexplained wage volatility—the variance of the permanent and transitory component of shocks to firm growth—may be invalid if the worker can influence the outcome of the firm. This could happen with the top managers of the firm because they exert a dominant role. To account for the possible bias induced by workers with dominant position inside the firm we focus on large firms, where arguably influence of any worker on firm productivity is diluted.

Our instruments may also be invalid if workers concentrate their stock investment in their firm's shares. This would give rise to an omitted variable problem because the portfolio share of risky asset is inversely related to the variance of risky asset returns (as in classical Merton-type portfolio choice models), which for investors holding significant shares of their firm may

be directly related to the variance of firm value added.²⁶ To account for potential instrument invalidity due to “own-firm bias” in household portfolio, we drop individuals with any holdings in their own firm.²⁷

Finally, a possible concern is that for a family what matters is the variation in total household earnings, rather than that of the primary earner. Indeed, within-family insurance (*e.g.* through added worker effects) may invalidate the use of the primary earner’s wage volatility as a measure of uninsured wage risk. To address this issue, we construct a measure of volatility based on household earnings (while continuing to use the same set of instruments as in the baseline regression—which refer to the primary earner). Similarly, one may be concerned that some individuals may try to smooth some of their wage risk by holding multiple concurrent jobs. To tackle this issue, we redefine our measure of wage volatility by summing wages from all jobs (rather than using just the main one). Correspondingly, we redefine our firm-specific instruments by weighting firm residuals by the share of total wages received by each firm.

Results for these various robustness checks are shown in Table 6. In column (1), we reproduce the baseline results. In columns (2) and (3), we report regressions when we retain only “large” firms (with size above the 25th and 50th percentile of the firm size distribution, respectively). As can be seen, these exclusions—if anything—strengthen the estimated marginal effect of background risk and leave our qualitative conclusions unchanged. In column (4), we drop workers who have some assets invested in their own firm. Since this group is rather small, the results barely change. Finally, in columns (5) and (6) we take broader measures of uninsurable wage risk, which include earnings of other household members²⁸ or earnings from secondary jobs, respectively. The results are again qualitatively unaffected. In both cases, instruments are less powerful but still pass conventional acceptability thresholds (and the Hansen test statistics reveal no sign of misspecification).

Sample selection: In our last set of robustness checks, we look at sample selection. Recall that our baseline sample is composed of individuals who are continuously employed in the private sector. This is because our instrument (firm shocks to value added) is undefined for those out of work and for those employed in the public sector. However, we can assess the robustness of our estimates by including in the sample individuals with spells in both the public and private sector (so that our measure of wage variability (7) changes because it includes past periods in the public sector), but continue to run our baseline regression (1) using only the periods in which individuals are observed working in the private sector (so that the firm-related instrument remains defined). Similarly, we can add to our sample those who had spells in the private sector amid spells of unemployment (resulting in yearly earnings less than the threshold amount for unemployment benefits, with similar adjustments to the variables). Results for these alternative samples are shown in Table 7. Clearly, we are now dealing with much larger samples. The estimates are slightly reduced, but again remain in the same ballpark.

In sum, looking at the broad picture coming out of Tables 5–7, the main message is one of remarkable stability of our main findings.

26. Døskeland and Hvide (2011) find that among Norwegian direct stockholders, 20% of the stock portfolio is held in shares of current or previous (last 10 years) employers.

27. The results are also robust to, instead of dropping individuals with holdings in their employers firm, redefining the risky portfolio to include only stocks in firms other than their own (*i.e.* the share of risky assets is redefined as $S'_{it} = \frac{R'_{it}}{R'_{it} + RF_{it}}$, with R' being risky assets net of the value of own-firm stocks).

28. Household earnings volatility is obtained using the same methodology described in Section 5.1 (*i.e.* the variance of the residual of a regression of household earnings on observables).

TABLE 6

Robustness: excluding small firms and ownership in own firm, including household and secondary jobs earnings

	Fixed effect IV (<i>Baseline</i>) and					
	Fixed effect IV (<i>Baseline</i>)	Restricting firm size: >25th perc. >50th perc.		Excluding owners	Obtaining σ_{it}^2 using: HH earnings second job	
	(1)	(2)	(3)	(4)	(5)	(6)
σ_{it}^2	-0.4986*** (0.1827)	-0.6262*** (0.2107)	-0.7558*** (0.2304)	-0.5161*** (0.1838)	-0.5665*** (0.2096)	-0.6116*** (0.2246)
Firm bankrupt in 1 year	-0.0201*** (0.0066)	-0.0207*** (0.0067)	-0.0208*** (0.0071)	-0.0201*** (0.0067)	-0.0224*** (0.0069)	-0.0204*** (0.0067)
Firm bankrupt in 3 years	-0.0040 (0.0034)	-0.0034 (0.0035)	-0.0035 (0.0036)	-0.0040 (0.0034)	-0.0043 (0.0034)	-0.0036 (0.0034)
Firm bankrupt in 5 years	-0.0020 (0.0035)	-0.0037 (0.0036)	-0.0046 (0.0037)	-0.0019 (0.0035)	-0.0019 (0.0035)	-0.0015 (0.0035)
Lagged log wealth	0.0112*** (0.0003)	0.0106*** (0.0003)	0.0102*** (0.0003)	0.0111*** (0.0003)	0.0115*** (0.0003)	0.0112*** (0.0003)
Home ownership	0.0146*** (0.0012)	0.0139*** (0.0013)	0.0122*** (0.0014)	0.0146*** (0.0012)	0.0146*** (0.0013)	0.0146*** (0.0012)
Age	0.0222*** (0.0017)	0.0215*** (0.0020)	0.0200*** (0.0022)	0.0220*** (0.0018)	0.0234*** (0.0014)	0.0222*** (0.0018)
Age sq.	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)
Family size	-0.0013* (0.0008)	-0.0014* (0.0008)	-0.0015* (0.0008)	-0.0012 (0.0008)	-0.0006 (0.0008)	-0.0013* (0.0008)
Hansen <i>J</i> -test <i>p</i> -value	0.13	0.19	0.36	0.12	0.14	0.15
<i>F</i> -stat first stage	56.85	44.53	38.04	56.09	35.63	44.07
Observations	1,184,800	1,124,682	1,038,205	1,173,031	1,184,800	1,184,800

Notes: The table reports estimates of the marginal effect of wage and unemployment risk on the risky financial portfolio share. Column (1) reproduces the IV benchmark regression of Table 3, column (3); columns (2) and (3) run the IV estimates on the sample of large firms, respectively, above the 25th percentile (column (2)) and the median size (column (3)). Column (4) excludes from the sample workers with own-firm stocks; column (5) measures wage risk with the variance of family earnings; columns (6) measures wage risk including also a worker's secondary job in a given year. Regressions also control for family type, area, and year dummies. Hansen *J*-test for instrument validity and *F*-stat for the power of the instruments are shown at the bottom of the table. Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

6.2.1. Heterogeneity. The effect of uninsurable wage risk on the demand for risky assets should be less important for households that have greater access to self-insurance (through buffers of accumulated assets). Similarly, pass-through coefficients of firm risk onto wages should be larger for wealthier individuals, as they are more willing to bear risk coming from the firm side due to their presumably higher risk tolerance.

These response heterogeneity predictions can be easily tested using interactions with household wealth. The results are reported in Table 8. In the top panel, we report pass-through estimates. The first two columns replicate the estimates of the model of Table 2 using our sample (instead of the universe of private sector workers). Ignoring interactions with wealth, pass-through estimates are (reassuringly) very similar to those reported in Table 2. The last two columns show pass-through estimates when permanent and transitory firm shocks are interacted with wealth. As expected, firms offer less insurance to workers with higher wealth, particularly against permanent shocks (the interaction with transitory shocks is not statistically significant).

In Panel B, we augment our baseline risky portfolio share regressions by interacting the variance of the worker's wage growth with lagged log financial wealth (and using as additional instruments the interaction of the latter with the firm's transitory and permanent shock variances).

TABLE 7
Robustness: including public sector and unemployment spells

	<i>Baseline including:</i>		
	Fixed effect IV (<i>Baseline</i>)	Public sector spells	Unemployment spells
	(1)	(2)	(3)
σ_{ii}^2	-0.4986*** (0.1827)	-0.4754*** (0.1405)	-0.3985*** (0.1199)
Firm bankrupt in 1 year	-0.0201*** (0.0066)	-0.0114** (0.0056)	-0.0122** (0.0052)
Firm bankrupt in 3 years	-0.0040 (0.0034)	0.0011 (0.0030)	0.0018 (0.0028)
Firm bankrupt in 5 years	-0.0020 (0.0035)	-0.0030 (0.0030)	-0.0019 (0.0035)
Lagged log wealth	0.0112*** (0.0003)	0.0107*** (0.0002)	0.0103*** (0.0002)
Home ownership	0.0146*** (0.0012)	0.0136*** (0.0010)	0.0136*** (0.0010)
Age	0.0222*** (0.0017)	0.0188*** (0.0018)	0.0171*** (0.0022)
Age sq.	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0001*** (0.0000)
Family size	-0.0013* (0.0008)	-0.0002 (0.0006)	0.0001 (0.0006)
Hansen <i>J</i> -test <i>p</i> -value	0.13	0.29	0.21
<i>F</i> -stat first stage	56.85	83.18	87.79
Observations	1,184,800	1,673,220	1,778,650

Notes: The table reports estimates of the marginal effect of wage and unemployment risk on the risky financial portfolio share. Column (1) reproduces the IV benchmark regression of Table 3; column (2) includes workers that during the sample period also have a spell in the public sector; column (3) includes workers that in a year receives a payroll of less than the social security basic amount, effectively including unemployed workers. Regressions also control for family type, area, and year dummies. Hansen *J*-test for instrument validity and *F*-stat for the power of the instruments are shown at the bottom of the table. Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

We find again intuitive results: the marginal effect of uninsurable wage risk on the demand for risky assets declines as the level of financial wealth increases.²⁹

29. These results can also be used to address the criticism that our estimate of the marginal effect of wage risk is high due to Local Average Treatment Effects (LATEs) (Angrist and Imbens, 1994). It is well known that in the presence of response heterogeneity, the IV estimator estimates (under some assumptions) not the “average treatment effect” (in our case, the average decline in the share of risky assets in portfolio that follows an increase in wage risk), but a “local average treatment effect”, which may be interpreted as the average treatment effect for the individuals who are mostly affected by a change in the instrument (*i.e.* the firm-related risk). For the Local Average Treatment Effect (LATE) interpretation to be responsible for the high value of our baseline estimate, we need the coefficient of the interaction in the pass-through regressions to be of opposite sign to the coefficient of the interaction in the share regressions (those mostly affected by the change in the instruments, *i.e.* those with a larger pass-through coefficient, should be the ones with the larger sensitivity of uninsurable wage risk to the demand for risky assets). However, we find exactly the opposite, suggesting that LATE is unlikely to be an issue. In unreported regressions (available on request), we generalize this exercise by allowing the partial insurance coefficients to vary with a whole vector of observable individual and firm characteristics: length and type of education, wealth, firm size, age, gender. And the same we do for the portfolio share equation. Though we find that some of these variables (namely schooling, wealth, and firm size) are significant shifters of the pass-through and/or of the effect of background risk on the share of risky assets in portfolio, we do not find anything systematic that would make us conclude that a LATE interpretation is justified.

TABLE 8
Wealth-induced heterogeneity

Panel A: Pass-through regressions				
	Permanent value added shocks	Transitory value added shocks	Permanent value added shocks	Transitory value added shocks
	(1)	(2)	(3)	(4)
Pass-through	0.0796*** (0.0053)	0.0153*** (0.0062)	0.0232 (0.0157)	0.0010 (0.0116)
Pass-through*lagged log wealth			0.0048*** (0.0012)	0.0012 (0.0010)
Hansen <i>J</i> -test <i>p</i> -value	0.62	0.00	0.63	0.00
<i>F</i> -stat first stage	178.91	752.05	107.26	172.64
Observations	1,316,004	1,321,303	1,316,004	1,321,303
Panel B: Risky share regressions				
σ_{it}^2			-2.1392*** (0.7799)	
σ_{it}^2 *lagged log wealth			0.1379** (0.0652)	
Hansen <i>J</i> -test <i>p</i> -value			0.27	
<i>F</i> -stat first stage			28.82	
Observations			1,184,800	

Notes: The table reports estimates of the pass-through coefficients (Panel A) and of the marginal effect of wage risk on the risky portfolio share (Panel B) allowing both coefficients to vary with the lagged value of household wealth. Regressions in Panel B also control for family type, area, and year dummies. Hansen *J*-test for instrument validity and *F*-stat for the power of the instruments are shown at the bottom of the table. Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

Figure 3 plots the pass-through effect (the dotted line on the left-hand scale, obtained considering permanent firm shocks only) and the marginal effect of uninsurable wage risk on the portfolio share for households (the continuous line on the right-hand scale) at different points of the distribution of wealth. Pass-through is always positive and it varies between 0.05 and 0.1 as wealth moves from the bottom to the top percentile.

The marginal effect of uninsurable wage risk on portfolio allocation is negative at all levels of wealth. However, while at the bottom of the distribution is large (around -1 or less), it is around -0.5 at median wealth and very close to zero at the top percentile of financial wealth—consistent with the prediction of a self-insurance model. As we discuss in the next session, this wealth-induced heterogeneity in workers' insulation from firms shocks and in response to wage risk translates in heterogeneity in the relevance of wage risk. Furthermore, since total wealth and even more so the holdings of risky assets are heavily concentrated, the effect of uninsurable wage risk on the aggregate demand for risky assets is likely small—a calculation we perform formally in the next section.

7. QUANTIFYING THE EFFECT OF UNINSURABLE WAGE RISK

The quantitative assessment of the importance of uninsurable wage risk hinges on two ingredients. The first is the size of λ , the marginal effect of a unit increase in the risk arising from on-the-job wage variation. A broad reading of our findings suggests that $\lambda \approx -0.5$ for the average/median wealth investor. We will use this value to perform our calculations. The second ingredient is the size of overall uninsurable wage risk. Gauging the latter is more problematic. We cannot use the

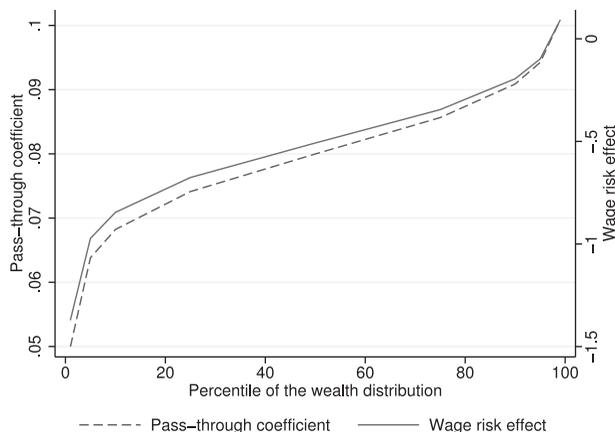


FIGURE 3

Wealth-induced heterogeneity in pass-through and marginal effect of wage risk

Notes: The dotted line shows the pass-through coefficient of permanent shocks to firm value added onto workers' wages, by worker wealth percentile (left-hand scale). The line is obtained using the estimates reported in Table 8, Panel A. The solid line shows the IV estimate of the marginal effect of wage risk on the share in risky assets, again by wealth percentile (right-hand scale). The line is obtained using the estimates of Table 8, Panel B.

size of the residual wage variance precisely because of the argument that not all residual variation is risk. However, we can provide bounds of the overall effect on the portfolio share.

From equation (3), uninsurable wage risk is defined as:

$$B_{it} = \theta_v^2 V_{it} + \theta_f^2 F_{it}$$

For given values of the estimated F_{it} and V_{it} —the variance of the firm's value added growth and the variance of the worker's earnings growth, respectively—the size of uninsurable wage risk depends on θ_v , the extent of worker-specific variation that is due to risk rather than choice, and the pass-through of firms shocks to wages θ_f . To assess the importance of uninsurable wage risk we consider two exercises. First, we compute the contribution of estimated wage risk to the portfolio share as:

$$\widehat{\lambda} \widehat{B}_{it} = \widehat{\lambda} \left(\widehat{\theta}_v^2 \widehat{V}_{it} + \widehat{\theta}_f^2 \widehat{F}_{it} \right) \quad (11)$$

Secondly, we estimate the effect on the risky portfolio share of changing wage risk from this estimated baseline by varying workers exposure to firm specific risk θ_f or increasing the share of worker-specific wage variation that is risk, θ_v :

$$\widehat{\lambda} \Delta B_{it} = \widehat{\lambda} \left((\theta_v^2 - \widehat{\theta}_v^2) \widehat{V}_{it} + (\theta_f^2 - \widehat{\theta}_f^2) \widehat{F}_{it} \right)$$

This computation assesses the economic importance of uninsurable wage risk by counterfactually “shocking” the two parameters that capture workers' exposure to risk, one through institutions or extent of superior information workers may have about evolution of their wages, θ_v ; the other through firm-provided insurance, θ_f . Shocking θ_f is of interest because, as shown by Lemieux *et al.* (2009) and Bénabou and Tirole (2016), there is strong evidence of a rise of pay for performance wage schemes and high-powered incentives over the past decade, not only

among workers in top positions but also among low-rank employees.³⁰ And competitive pressure for talent could make incentives even more powered in the future.

To perform these calculations, we take the pass-through coefficient with respect to permanent firm shocks, $\widehat{\theta}_f = 0.07$ (because the response to transitory shocks is tiny, and hence adding it would make little difference). To quantify the baseline share of worker-specific wage variation that is risk, $\widehat{\theta}_v$, we consider the following argument: under the assumptions that censoring bias is unimportant, insurance within the firm is substantial (both backed by the estimates in Table 2 and the evidence in Table 4), and measurement error is negligible (given the administrative nature of the data), then $p\lim \widehat{\lambda}_{FE} \approx \widehat{\theta}_v^2 p\lim \widehat{\lambda}_{IVFE}$. Hence, one can infer that $\widehat{\theta}_v \approx 0.2$. We might expect $\widehat{\theta}_v$ to represent a lower bound for the true θ_v , both because of some measurement error in wages and because there may be additional sources of background risk that our empirical strategy is missing. Finally, we estimate F_{it} and V_{it} using the variance of the firm's value added growth and the variance of the worker's earnings growth, respectively: $\widehat{F}_{it} = 0.16$ and $\widehat{V}_{it} = 0.053$, from Table 1.³¹

The surface we plot in Figure 4 is the economic effect of uninsurable wage risk on the share of risky assets in portfolio, computed using equation (11) as:

$$\widehat{\lambda}(\theta_v^2 \widehat{V}_{it} + \theta_f^2 \widehat{F}_{it})$$

where we use the baseline estimate $\widehat{\lambda} = -0.5$. The crossing between the two darker lines on the surface marks the sample estimates combination $(\widehat{\theta}_v, \widehat{\theta}_f)$.

Evaluated at the average values of V_{it} and F_{it} and at the point estimates of the parameters $(\widehat{\lambda}, \widehat{\theta}_v, \widehat{\theta}_f)$ the economic effect of uninsurable wage risk is tiny: the predicted decline in the share of risky assets is -0.14 percentage points. However, if workers were to share equally the firm-specific risk ($\theta_f = 0.5$), for given θ_v , the effect would be as high as 2 percentage points (or 10% of the average share of risky assets in portfolio). In contrast, holding constant θ_f , while increasing the amount of worker-specific variation that is due to risk rather than choice, leaves the effect of background risk on the demand for stocks fairly small. Indeed, even if half of the worker-specific wage variation was risk, the effect of uninsurable wage risk would remain small: a predicted 0.7 percentage point decline. This is visible from the slope of the surface, which is steeper when we move along the θ_f -axis than when we move along the θ_v -axis. This difference in the effect of an equally sized increase in θ_f and θ_v follows from the larger variance of firms shocks \widehat{F}_{it} , which is 3 times larger than the worker specific wage variation \widehat{V}_{it} , implying a much larger increase in wage risk when θ_f rather than θ_v increases.

We have documented substantial wealth-induced heterogeneity in the pass-through of firm-related shocks onto wages as well as in the sensitivity of the demand for stocks to uninsurable wage risk. Consequently, we should expect substantial heterogeneity in the economic effect of the latter. To illustrate, we consider the effect for households at the 5th and 95th percentile of the wealth distribution. The estimates of $\widehat{\lambda}$ are, respectively, -0.97 and -0.097 . The other important element that varies is the pass-through coefficient, which takes values 0.06 and 0.10, respectively, for the 5th and 95th percentile of the wealth distribution. Evaluated at the average values of V_{it}

30. Lemieux *et al.* (2009) show that in the U.S. between the 1970s and the 1990s, the fraction of workers paid on the basis of performance rose from 38% to 45%, and for salaried workers from 45% to 60%. This pattern is not confined to the U.S. Bloom and Van Reenen (2011), for instance, document that the fraction of U.K. establishments using some form of performance pay rose from 41% in 1984 to 55% in 2004.

31. In fact, an estimate of V_{it} should subtract, from the variance of wage growth, the contribution of the firm component—which is, however, tiny given the extent of insurance within the firm.

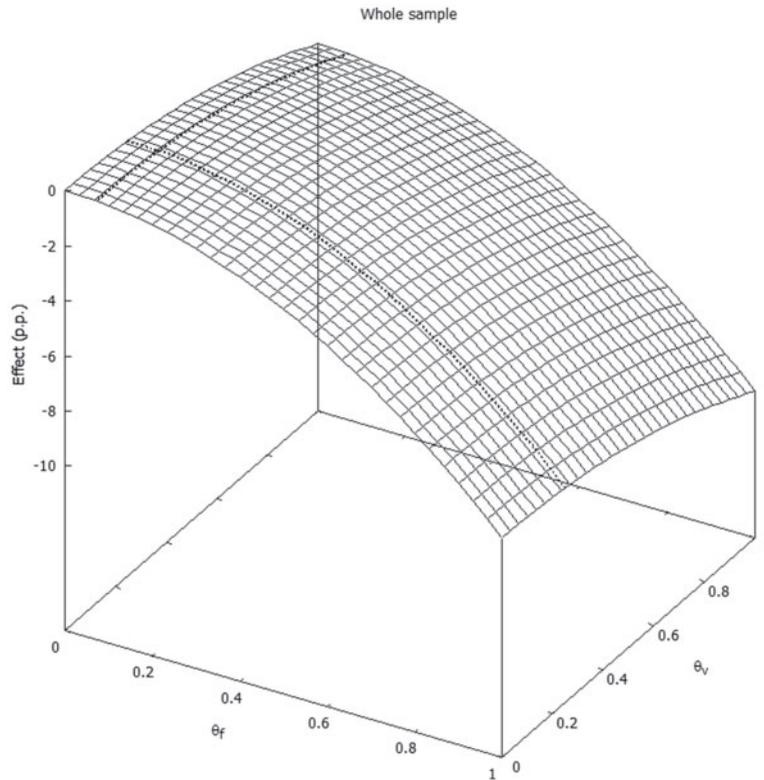


FIGURE 4

The effect of wage risk on the share of risky assets

Notes: The surface plot represents the percentage points reduction (on the vertical line) in the portfolio share in risky assets by Norwegian households as we vary the level of insurance within the firm (θ_f) and the fraction of non-firm related wage variation that is not avoidable (θ_v). The crossing of the darker lines on the surface corresponds to the point estimates of the pass-through of permanent firm shocks on wages and the share of non-firm related wage variability that is not avoidable (from the baseline specifications).

and F_{it} and at the point estimates of the parameters $\hat{\theta}_v, \hat{\theta}_f$, the economic effect of background risk are still small in both groups (-0.23 percentage points at the 5th wealth percentile and -0.06 percentage points at the 95th percentile). Figure 5 reports the corresponding wage risk effect surfaces for the two groups.

For the wealthy, neither variations in θ_f nor θ_v would affect their background risk response much. The response surface is fundamentally flat. In contrast, the slope of the surface among the poor is much steeper; a reduction in firm insurance could potentially have large impact on their portfolio choice, reducing even further the amounts of wealth held in risky instruments. For these workers, sharing half of the shocks to their firms would lower the portfolio share in risky assets by about 15 percentage points, a very large drop. Also an increase in wage risk unrelated to the firm's fortunes could have a substantial impact. However, because these workers own a small fraction of total stocks, these larger effects are unlikely to generate large aggregate consequences (which we document next).

As our last exercise, we look at the effect of uninsurable wage risk for the aggregate demand for stocks in the baseline and in the hypothetical scenarios in which we vary the extent of risk faced by individuals. We allow for wealth-related heterogeneity in both the pass-through of firms shocks

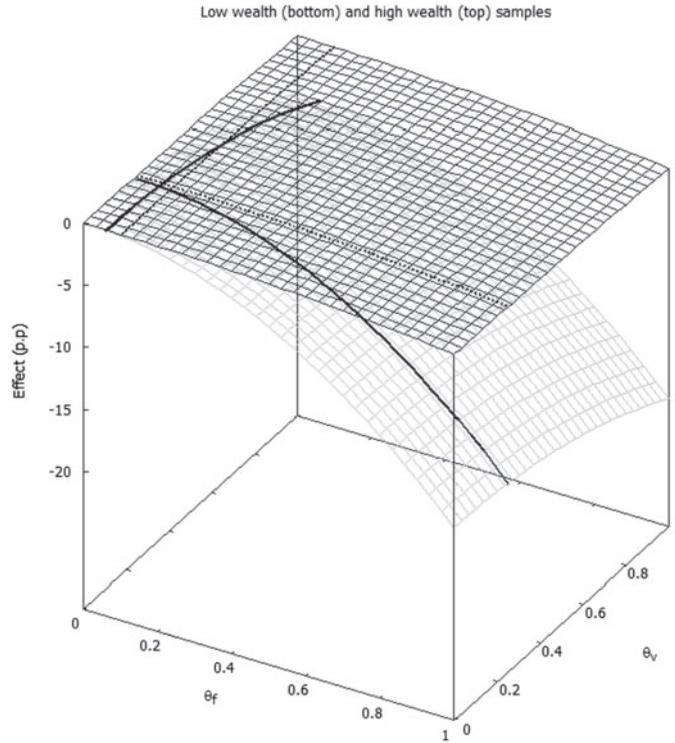


FIGURE 5

The effect of wage risk for low- and high-wealth samples

Notes: The surface plots represent the percentage points reduction (on the vertical line) in the portfolio share in risky assets held by rich households (corresponding to the 95th percentile of the wealth distribution, the top surface) and poor households (corresponding to the 5th percentile of the wealth distribution, the bottom surface). Movements along the surfaces are due to different combinations of the level of insurance within the firm (θ_f) and the fraction of non-firm related wage variation that is not avoidable (θ_v). The crossing of the darker lines on the two surfaces corresponds to the point estimates of the pass-through of permanent firm shocks to wages and the share of wage variability that is not avoidable (separately by wealth, see Table 8).

and the portfolio sensitivity to uninsurable wage risk. This exercise is relevant for understanding the role of uninsurable wage risk for asset prices.

To perform this exercise, we consider an increase in θ_f and θ_v from their point estimate to 0.5, so that workers share 50% of the permanent shocks to their firm and 50% of their personal wage variation is risk. For a given worker i with initial wealth A_{it-1} , the effect on the risky share of raising θ_v and θ_f from $(\hat{\theta}_v, \hat{\theta}_f)$ to (0.5, 0.5) is:

$$\Delta S_i = (\hat{\lambda}(A_{it-1})(0.25\hat{V}_{it} + 0.25\hat{F}_{it}) - \hat{\lambda}(A_{it-1})(\hat{\theta}_v^2\hat{V}_{it} + (\hat{\theta}_f(A_{it-1}))^2\hat{F}_{it}))$$

and that on the individual demand for stocks:

$$\text{Change in demand for stocks} = A_{it-1} \Delta S_{it}$$

Accordingly, our estimate of the effect on the aggregate demand for stocks is

$$\% \text{change in aggregate demand for stocks} = \left(\sum_i A_{it-1} \Delta S_{it} \right) / (\text{Total stocks}_{t-1})$$

We estimate this effect to be 0.2% on average over all sample years—a tiny response to a large change in uninsurable wage risk. Increasing the size of the shock by setting θ_f and θ_v to 0.8 leaves the result qualitatively unchanged. The reason why the aggregate demand for stocks is insensitive to wage risk is that the effect is small at high wealth levels, and the ownership of risky assets is concentrated precisely among the wealthy (the top 5% holds 78% of the stock market in Norway and the top 10% holds 87%). In fact, we calculate that among the households with below median wealth increasing θ_f and θ_v to 0.5 lowers the demand for risky assets by 2.8% while it has a negligible effect among households with above median wealth.

Overall, the calculations in this section imply that uninsurable wage risk is economically important for individuals with low assets; for those who can count on a sufficiently high level of buffer savings the tempering effect of uninsurable wage risk is contained. The combination of very high sensitivity among the poor, low sensitivity among the wealthy, and the concentration of risky assets in the hands of the latter implies a small effect of even large increases in uninsurable wage risk on the aggregate demand for risky assets. This suggests a small role of uninsurable wage risk as a driver of asset prices.

8. CONCLUSIONS

In this article, we have argued that achieving identification of the effect of wage risk on people's financial decisions poses important conceptual challenges and formidable data requirements. Using extremely rich and virtually error-free Norwegian administrative data we estimate firm-related measures of workers' earnings variation to isolate exogenous changes in uninsurable wage risk. We show that once the endogeneity of usual measures of wage risk is properly addressed and unobserved heterogeneity and censoring of stock investments are accounted for, the estimated sensitivity of the risky portfolio share of the average household to earnings risk can be up to 25 times larger than the estimates obtained ignoring these issues, as done in the existing literature. Furthermore, the sensitivity differs depending on wealth buffers accumulated by the household, spanning from irrelevance at the top of the wealth distribution to sensitivity that is 3 times the average for people at the bottom. While sensitivity to uninsurable wage risk is very large, we find small sensitivity to unemployment risk, possibly because of generous UI in Norway.

Can uninsurable wage risk explain the large amount of heterogeneity in portfolio composition observed in the data? Answering this question requires a consistent estimate of the marginal effect of uninsurable wage risk, which we have, and a comprehensive measure of the size of uninsurable wage risk, which we can bound. At sample means and for the median wealth household the contribution of uninsurable wage risk is small, not because sensitivity is small but because firms provide substantial wage insurance to their workers.³² However, because marginal responses differ considerably depending on the buffers accumulated, the economic importance of uninsurable wage risk varies greatly: it is relatively large for the poor and negligible for the wealthy. In this sense, uninsurable wage risk is a viable explanation of portfolio heterogeneity among low-wealth people but not among the high-wealth segment. Moreover, as the wealthy hold most of the risky assets in the economy, wage risk is unlikely to matter for asset pricing. Finally, our empirical strategy identifies the parameters that drive uninsurable wage risk and the portfolio response to it, hence allowing us to run meaningful counterfactuals. We show that the adoption of high-powered wage incentive contracts can cause large responses in portfolio

32. This result is in line with Berk and Walden (2013). In their paper, the authors propose that most workers get their optimal risk allocation through the employment contract, and, therefore, have no need to participate in stock markets. It is thus not that agents do not care about background risk, but rather that their background risk is well aligned with their consumption risk. Our identification strategy separates scope from motive and identify the latter.

composition, particularly for people below median wealth. In sum, our strategy allows to say when, for what and for whom wage risk matters: it matters when firms offer little insurance, it matters for understanding why people with low-wealth buffers are reluctant to invest in stocks, and it is unlikely to matter for stock prices.

The idea of separating *motive* from *scope* that guides our methodology to assess the role of wage risk for the demand of risky assets can be extended to study portfolio hedging against human capital risk. If the firm component of workers' wages correlates differentially with single listed stocks, then workers may rely on the stock market to hedge their human capital risk, as first noted by Myers (1972). However, if workers receive considerable wage insurance from their firm the scope for hedging can be limited even though the motive may be strong. Because we can identify the degree of insurance (and note that this is partial), our methodology can be used to identify the strength of the hedging motive. The pursuit of this idea is the next step in our research agenda.

In this article, we have focused on two sources of background risk—wage risk and unemployment risk. Given the large weight that human wealth has in the lifetime resources of most individuals, these are probably the most important sources of background risk for most individuals. But they are not the only ones. For homeowners, unanticipated shocks to housing wealth is another, and given the illiquidity of housing it cannot easily be avoided; for entrepreneurs, private business wealth, is still another—and has been studied by Heaton and Lucas (2000a, 2000b). These different sources of background risk share one common feature: each one accounts for a substantial share of a consumer lifetime resources. Thus, even if the effect of each one may be relatively contained, their joint effect on households assets allocation may be substantial. We have contributed to quantify one of them. More work is needed to quantify the others.³³

APPENDIX A

A.1. Data sets

The analysis uses several data sources maintained by Statistics Norway that can be combined through unique personal and household identifiers over time.

A.1.1. The Central Population Register. The Central Population register contains end of year information on all Norwegian residents for the time period 1993–2011 and contains individual demographic information (*i.e.* gender, day of birth, county of residence, and marital status). It also contains family identifiers allowing us to match spouses and cohabiting couples with common children. Identifying unmarried couples without common children is not possible in our sample period.

A.1.2. Administrative Tax and Income Records. Because households in Norway are subject to a wealth tax, they are every year required to report their complete income and wealth holdings to the tax authority, and the data are available every year from 1993 to 2011. Each year, before taxes are filed in April (for the previous year), employers, banks, brokers, insurance companies, and any other financial intermediaries are obliged to send both to the individual and to the tax authority, information on the value of the asset owned by the individual and administered by the employer or the intermediary, as well as information on the income earned on these assets. In case an individual holds no stocks, the tax authority pre-fills a tax form and sends it to the individual for approval; if the individual does not respond, the tax authority considers the information it has gathered as approved. In 2011, as many as 2.4 million individuals in Norway

33. Palia *et al.* (2014) study the effect of volatility in returns to human capital, housing and private equity on the risky portfolio share. Unfortunately, their study suffers from the endogeneity issues that we have stressed in this study (as it assumes that all measured variation in labour income, housing, and private equity returns is background risk). Calibration exercises show the potential importance of housing return risk for the composition of the financial portfolio (Cocco *et al.*, 2005) and of returns to private wealth (Heaton and Lucas, 2000b). But a proper empirical assessment of these sources is still missing and faces the same identification problems as those faced by human capital risk.

(66% of the tax payers) belonged to this category.³⁴ If the individual or household owns stocks, then he has to fill in the tax statement—including calculations of capital gains/losses and deduction claims. The statement is sent back to the tax authority, which, as in the previous case receives all the basic information from employers and intermediaries and can thus check its truthfulness and correctness. Stockholders are treated differently because the government wants to save on the time necessary to fill in more complex tax statements and to reduce the risk of litigation due to miscalculated deductions on capital losses and taxes on capital gains. Traded financial assets are reported at market value. For stocks in non-listed companies that are not traded the company itself has to provide a tax report to the tax registry every year. In this report, the company proposes a value of the company by the end of the year. This value should be the total net worth of the company, after deducting any debts. All assets have to be included in the valuation, except goodwill which is not included. The tax authority may adjust the value of the company upwards after going over the report, if it does not find the proposed value reasonable. Obviously this leads to undervaluation of the companies, but this is bound as unrealistically low figures would cause the tax authority to start a more thorough investigation.

This procedure, particularly the fact that financial institutions supply information on their customers' financial assets directly to the tax authority, makes tax evasion very difficult, and thus non-reporting or under-reporting of assets holdings are likely to be negligible.

A.1.3. The Norwegian National Educational Database. Educational attainment is reported by the educational establishment directly to Statistics Norway at the individual level, hence minimizing the measurement error. The information includes on every student the highest level of education) at the individual level as of October every year.

A.1.4. The Register of Shareholders. The register contains ownership of all Norwegian limited liability companies (related to the statistics on "Aksjer og kapitalutdelinger"). Importantly, the register contains information about shareholders and received dividends. Dividends are reported at the yearly level, and ownership is reported as of 31 December each year.

A.1.5. Employer–Employee Register. All firms hiring workers in Norway are required to report all work relationships to the Central Employer–Employee register. This includes registering the date and individual ID for the each time an employment relationship is established or terminated and when permanent changes are made to the registered information about working hours, job title (occupation code), and workplace (department). The register also contains the organization number of the firm and the sum of total payments (wages and remuneration) from the firm to the worker at a yearly level. When a worker has work relationships with several firms during the year, we select the firm with the highest payments to the worker that year as the main work relationship.

A.1.6. The Central Register of Establishments and Enterprises. The register contains all enterprises and establishments in the private and public sector in Norway. For our purposes, we select information on organization ID, geographical information, institutional sector, industrial classification (NACE), number of employees.

A.1.7. Firm Balance Sheet Register. The register contains accounts and balance sheet information from the financial statements of all non-financial firm. We extract all variables needed to calculate value added per worker. Some of the main variables and definitions are as follows:

Operating income and operating expenses are ordinary income and expenses outside financial ones. Operating income is divided into sales revenues (taxable and tax-free), rental income, commission revenues, profits from the sale of fixed assets, and other operating-related revenues. Operating expenses include changes in stocks, costs of raw materials and consumables used, wages and salaries, depreciation and write-downs of tangible fixed assets and intangible fixed assets as well as a number of different types of other operating expenses. Examples of operating expenses that are specified are subcontracting, repair and maintenance, and expenses relating to means of transport.

Cost of raw materials and consumables used includes stock changes of work in progress and finished goods.

Wages and salaries include wages, holiday pay, employers' national insurance premium, pension costs, and other personnel expenses.

Financial income and financial expenses are ordinary revenues and expenses relating to investments, securities, receivables, and liabilities. The financial items also include share of earnings relating to foreign exchange gains and losses (*agio*) and value changes of market-based current asset investments.

Extraordinary revenues and expenses apply to material items that are unusual for the business and do not occur regularly.

34. See the 2011 Annual Report from the Norwegian Tax Administration, <http://www.skatteetaten.no/en/>.

Taxes represent taxes relating to the accounting result, and consist of taxes payable, expected reimbursement claims from owners, and changes in deferred taxes. Taxes payable are the taxes expected to be assessed on the year's taxable income corrected for any discrepancy between calculated and assessed taxes the year before.

Allocation of the profit/loss for the year shows how a profit is allocated and losses are covered. It provides information on transfers to/from equity and dividends to owners.

Fixed assets cover assets that are mainly included in the enterprise's long-term creation of value and are intended for permanent ownership or use, as well as receivables and securities scheduled for repayment later than 1 year after the time of settlement. This includes tangible fixed assets broken down into buildings and facilities, facilities under construction, transport equipment, machinery, etc. Long-term receivables and investments are included as fixed assets, such as investments in other activities and loans to enterprises in the same group.

Current assets are assets relating to the enterprise's sales of goods and services, or which are expected to have a functional period of less than 1 year in operation. This includes cash and short-term capital investments (cash, bank deposits, shares, bonds, etc.), receivables and inventories. Receivables are current assets if it has been agreed or scheduled that they shall be repaid within 1 year after the end of the financial year.

Equity is the portion of the total capital belonging to the owners, and is shown as the value of assets less liabilities. Equity is classified in two main divisions, invested equity and retained earnings. Invested equity consists of share capital and share premium accounts. Retained earnings consist of fund for assessment differences and other reserves/uncovered losses.

Liabilities cover all obligations that can come to place restrictions on the future use of the enterprise's resources, and are divided into provisions for liabilities and charges (pension commitments, deferred tax liabilities, etc., other long-term liabilities and short-term liabilities. Long-term liabilities are legal or financial obligations not meant to be redeemed during the coming accounting period, and are not related to the enterprise's short-term sales of goods and services. Short-term liabilities are liabilities that fall due for payment within 1 year from the time of settlement, or are directly related to the enterprise's short-term sales of goods and services.

A.1.8. Register of Bankruptcies. The register contains the firm number and the exact date of bankruptcy at the firm level. All juridical objects, which includes all types of firms/enterprises and individuals who have unpaid accounts and are by definition insolvent, can be declared bankrupt.

A.2. Sample selection

We start with a data set on income recipients that merges record from the Central Population Register and the Administrative Tax and Income Register. This merged data set includes 29,814,364 person-year observations for the period 1995–2010. Given that we need to use as an instrument a measure of firm-level risk, we focus on a sample of individuals who are continuously employed in the private sector (sector 710 or 717). This excludes those who are not working (unemployed, retired, disabled, etc.) and those who have a spell in the government sector. This sample selection leaves us with 9,888,562 observations. Next, we exclude individuals who are younger than 25 years (and hence possibly still in school) and those older than 60 years (who may have intermittent participation, and also have widespread access to early retirement, typically from the age of 62 years). We are left with 7,566,412 observations. Merging this data set with firm-level information reduces the usable sample to 6,501,730 observations (this sample reduction is due to some missing information in the firm data set used to construct the measure of firm value added, exclusion of short lived firms—those that are active for less than 3 years—and some inconsistencies in the reported firm number in the Employer/Employee Registry versus the Balance Sheet Registry). Next, we exclude individuals who have earnings below the basic amount threshold of the Norwegian Social Insurance Scheme (*grunnbeløpet*) in one or more years and are left with 5,168,462 observations. Even though we restrict the sample of workers between 25 and 60 years of age, some students are still left in the sample, and will typically have low incomes.³⁵ Further, workers who have some period of disability or sick leave, will often have less than full-time positions, potentially in several firms. To reduce the impact of such outliers, we drop all the observations where earnings growth is less than –80% or more than 500% (and are left with 5,115,196 observations). Since we run regressions at the household level, we keep only the primary earner of the household (4,846,766 observations left). The number of observations in the various regressions we run are less than this because we use lags for constructing some of the variables and instruments.

35. The incentive to stay below this threshold is significant as the government stipend to all students is reduced almost one-to-one for each dollar earned above a threshold only marginally higher than *grunnbeløpet*.

A.3. Mobility regressions

TABLE A1
Worker mobility

	Mover	Mover
	(1)	(2)
g_{jt}	0.242*** (0.0292)	0.242*** (0.0292)
g_{jt+1}	0.0278 (0.0267)	0.0278 (0.0267)
g_{jt+2}	0.0256 (0.0193)	0.0256 (0.0193)
Firm bankrupt in 1 year		0.453 (0.319)
Firm bankrupt in 2 years		0.0769 (0.344)
Observations	3,219,340	3,219,340

Notes: The table reports marginal effect estimates of a probit regression for the event of job mobility as a function of current (g_{jt}) and future firm shocks and worker's socio-demographic characteristics. In addition to the reported coefficients, the regressions also control for log firm size, dummy variables for individuals being recipients of sickness money, maternity/paternity benefits, UI, as well as education indicators (type and length), family type, area dummies, a quadratic in age, gender, and year dummies. Clustered standard errors are in brackets. Coefficient significance: *** at 1% or less; ** at 5%; * at 10%.

Acknowledgments. We thank Sumit Agarwal, Ådne Cappelen, Francis Vella, the editor, and three anonymous reviewers for helpful comments. We are grateful to seminar participants at Berkeley, Naples, Lugano, Geneva, Maastricht, Princeton, the 2014 SITE, the 2014 CEAR, the 2015 European Household Finance conference, and the 2017 ASSA Annual Meeting. An earlier version of this article circulated under the title "Back to Background Risk?". We are grateful to Finansmarkedsfondet (The Research Council of Norway, Grant #230843) for financial support and to Davide Malacrino for research assistance.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

REFERENCES

- AIYAGARI, R. S. (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics*, **109**, 659–84.
- ANGERER, X. and LAM, P.-S. (2009), "Income Risk and Portfolio Choice: An Empirical Study", *Journal of Finance*, **64**, 1037–1055.
- ANGRIST J. D. and IMBENS, G. (1994), "Identification and Estimation of Local Average Treatment Effects", *Econometrica*, **62**, 467–476.
- ARRONDEL, L. and CALVO-PARDO, H. (2012), "Endogenous Non-tradable Earnings Households' Demand for Risky Assets" (Discussion Papers in Economics and Econometrics No. 1414, University of Southampton).
- BÉNABOU, R. and TIROLE, J. (2016), "Bonus Culture: Competitive Pay, Screening, and Multitasking", *Journal of Political Economy*, **124**, 305–370.
- BERK, J. and WALDEN, J. (2013), "Limited Capital Market Participation and Human Capital Risk", *Review of Asset Pricing Studies*, **3**, 1–37.
- BETERMIER, S., JANSSEN, T., PARLOUR, C. A., *et al.* (2012), "Hedging Labor Income Risk", *Journal of Financial Economics*, **105**, 622–639.
- BISCOURP, P., DESSY, O. and FOURCADE, N. (2005), "Les salaires sont-ils rigides? Le cas de la France a la fin des années 90", *Economie et Statistique*, **386**, 59–79.
- BLOOM, N., VAN REENEN, J. (2011), "Human Resource Management and Productivity", in Card, D. and Ashenfelter, O., (eds), *Handbook of Labor Economics*, Vol. 4b (North Holland: Elsevier) 1697–1767.

- BLUNDELL, R. W. and SMITH, R. J. (1986), "An Exogeneity Test for a Simultaneous Equation Tobit Model with an Application to Labor Supply", *Econometrica*, 679–685.
- CALVET, L., CAMPBELL, J. and SODINI, P. (2009), "Measuring the Financial Sophistication of Households", *American Economic Review P&P*, **99**, 393–98.
- CAMPBELL, J. Y. (2006), "Household Finance", *The Journal of Finance*, **61**, 1553–1604.
- CARDOSO, A. R. and PORTELA, M. (2009), "Micro Foundations for Wage Flexibility: Wage Insurance at the Firm Level", *Scandinavian Journal of Economics*, **111**, 29–50.
- CHAMBERLAIN, G. (1984), "Panel Data", in Griliches, Z. and Intriligator, M., (eds), *Handbook of Econometrics*, Vol. 2 (North Holland: Elsevier) 1247–1318.
- COCCO, J. F. (2004), "Portfolio Choice in the Presence of Housing", *Review of Financial Studies*, **18**, 535–567.
- COCCO, J. F., GOMES, F. and MAENHOUT, P. (2005), "Consumption and Portfolio Choice over the Life-Cycle", *Review of Financial Studies*, **18**, 491–533.
- COCHRANE, J. H. (2006), "Financial Markets and the Real Economy", in Cochrane, J. H., (ed.) *International Library of Critical Writings in Financial Economics*, Vol. 18 (London: Edward Elgar). Available at: <https://www.elgar.com/shop/financial-markets-and-the-real-economy?>
- CONSTANTINIDES, G. and DUFFIE, D. (1996), "Asset Pricing with Heterogeneous Consumers", *Journal of Political Economy*, **104**, 219–240.
- CUNHA, F. and HECKMAN, J. J. (2007), "The Evolution of Inequality, Heterogeneity and Uncertainty in Labor Earnings in the U.S. Economy" (NBER Working Paper No. 13526).
- CUNHA, F., HECKMAN, J. J. and NAVARRO, S. (2005), "Separating Uncertainty from Heterogeneity in Life Cycle Earnings", *Oxford Economic Papers*, **57**, 191–261.
- DØSKELAND, T. and HVIDE, H. (2011), "Do Individual Investors Have Asymmetric Information Based on Work Experience?", *The Journal of Finance*, **66**, 1011–1045.
- EDWARDS, R. D. (2008), "Health Risk and Portfolio Choice", *Journal of Business and Economic Statistics*, **26**, 472–485.
- FRIEDRICH, B., LAUN, L., MEGHIR, C., et al. (2015), "Earnings Dynamics and Firm-Level Shocks" (Working Paper).
- FUSS, C. and WINTR, L. (2008), "Rigid Wages and Flexible Labour? Firm-level Evidence Based on Productivity for Belgium" (National Bank of Belgium Working Paper No. 159).
- GEORGARAKOS, D. and INDERST, R. (2014), "Financial Advice and Stock Market Participation" (CEPR Discussion Paper No. 9922).
- GOURINCHAS, P.-O. and PARKER, J. A. (2002), "Consumption over the Life Cycle", *Econometrica*, **70**, 47–89.
- GUERTZGEN, N. (2014), "Wage Insurance within German Firms: Do Institutions Matter?", *Journal of the Royal Statistical Society*, **177**, 345–369.
- GUIO, L., JAPPELLI, T. and TERLIZZESE, D. (1996), "Income risk, Borrowing Constraints and Portfolio Choice", *American Economic Review*, **86**, 158–172.
- GUIO, L., PISTAFERRI, L. and SCHIVARDI, F. (2005), "Insurance within the Firm", *Journal of Political Economy*, **113**, 1054–1087.
- GUIO, L. and SODINI, P. (2013), "Household Finance: An Emerging Field", *Handbook of the Economics of Finance* (North Holland, Amsterdam: Elsevier).
- GUVENEN, F. and SMITH, A. (2014), "Inferring Labor Income Risk and Partial Insurance from Economic Choice", *Econometrica*, **82**, 2085–2129.
- HEATON, J. and LUCAS, D. (1996), "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing", *Journal of Political Economy*, **104**, 443–487.
- HEATON, J. and LUCAS, D. (2000a), "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk", *Journal of Finance*, **55**, 1163–1198.
- HEATON, J. and LUCAS, D. (2000b), "Portfolio Choice in the Presence of Background Risk", *The Economic Journal*, **110**, 1–26.
- HEATON, J. and LUCAS, D. (2008), "Can Heterogeneity, Undiversified Risk, and Trading Frictions Solve the Equity Premium Puzzle?," in Mehra, R., (ed.), *Handbook of the Equity Risk Premium* (San Diego: Elsevier) 535–57.
- HOCHGUERTEL, S. (2003), "Precautionary Motives and Portfolio Decisions", *Journal of Applied Econometrics*, **18**, 61–77.
- HONORÉ B. E. and HU, L. (2004), "Estimation of Cross Sectional and Panel Data Censored Regression Models with Endogeneity", *Journal of Econometrics*, **122**, 293–316.
- HUNG, W., LIU, Y.-J., TSAI, C.-F., et al. (2014), "Employer Stock Risk, Employee Income Risks, and Portfolio Choice: New Evidence from Taiwan" (Mimeo, National Taiwan University).
- HURST, E., LI, G. and PUGSLEY, B. (2015), "Are Household Surveys Like Tax Forms: Evidence from Income Underreporting of the Self Employed", *Review of Economics and Statistics*, **96**, 19–33.
- JUHN, C., McCUE, K., MONTI, H. et al. (2017), "Firm Performance and the Volatility of Worker Earnings" (NBER Working Paper No.23102, January 2017).
- KATAY, G. (2009), "Do Firms Provide Wage Insurance against Shocks? Evidence from Hungary" (Working Paper No. 964, European Central Bank).
- KAUFMANN, K. and PISTAFERRI, L. (2009), "Disentangling Insurance and Information in Intertemporal Consumption Choices", *American Economic Review*, **99**, 387–392.
- KIMBALL, M. S. (1993), "Standard Risk Aversion", *Econometrica*, **61**, 589–611.
- LAGAKOS, D. and ORDONEZ, G. L. (2011), "Which Workers Get Insurance Within the Firm?", *Journal of Monetary Economics*, **58**, 632–645.

- LEMIEUX, T., MACLEOD, W. B. and PARENT, D. (2009), "Performance Pay and Wage Inequality", *Quarterly Journal of Economics*, **124**, 1–49.
- LEVHARI, D. and WEISS, Y. (1974), "The Effect of Risk on the Investment in Human Capital", *American Economic Review*, **64**, 950–963.
- LOKEN, E. and STOKKE, T. A. (2009), "Labor Relations in Norway" (Fafo Report No. 33).
- LOW, H., MEGHIR, C. and PISTAFERRI, L. (2010), "Wage Risk and Employment Risk over the Life Cycle", *American Economic Review*, **100**, 1432–1467.
- MEGHIR, C. and PISTAFERRI, L. (2004), "Income Variance Dynamics and Heterogeneity", *Econometrica*, **72**, 1–32.
- MERTON, R. (1971), "Optimum Consumption and Portfolio Rules in a Continuous-Time Model", *Journal of Economic Theory*, **3**, 373–413.
- MYERS, D. (1972), "Nonmarketable Assets and Capital Market Equilibrium under Uncertainty", in Jensen, M. C., (ed.), *Studies in the Theory of Capital Markets* (New York: Praeger) 223–48.
- NILSEN, Ø. A. and REISO, K. H. (2014), "Scarring Effects of Early-Career Unemployment", *Nordic Economics Policy Review*, **1**, 13–45.
- PALIA, D., QI, Y. and WU, Y. (2014), "Heterogeneous Background Risks and Portfolio Choice: Evidence from Micro-Level Data", *Journal of Money, Credit and Banking*, **46**, 1687–1720.
- PERESS, J. (2004), "Wealth, Information Acquisition and Portfolio Choice", *The Review of Financial Studies*, **17**, 879–914.
- POLKOVNICHENKO, V. (2007), "Life-cycle Portfolio Choice with Additive Habit Formation Preference and Uninsurable Labor Income Risk", *Review of Financial Studies*, **20**, 83–124.
- PRIMICERI, G. and VAN RENS, T. (2009), "Heterogeneous Life Cycle Profiles, Income Risk and Consumption Inequality", *Journal of Monetary Economics*, **56**, 20–39. Carnegie-Rochester Conference Series on Public Policy: The Causes and Consequences of Rising Economic Inequality.
- SCHMIDT, L. D. W. (2016) "Climbing and Falling Off the Ladder: "Asset Pricing Implications of Labor Market Event Risk". Available at SSRN: <http://dx.doi.org/10.2139/ssrn.2471342>.
- SMITH, R. J. and BLUNDELL, R. W. (1986), "An Exogeneity Test for a Simultaneous Equation Tobit Model with an Application to Labor Supply", *Econometrica*, **54**, 679–685.
- VICEIRA, L. (2001), "Optimal Portfolio Choice for Long-horizon Investors with Nontradable Labor Income", *The Journal of Finance*, **56**, 433–470.