

Firm-Related Risk and Precautionary Saving Response

By ANDREAS FAGERENG, LUIGI GUISO AND LUIGI PISTAFERRI*

We propose a new approach to identify the strength of the precautionary motive and the extent of self-insurance in response to earnings risk based on empirical Euler equation estimates. As is well known (see e.g. Carroll, 1992), households who face income risk and have a precautionary motive for saving accumulate wealth to cushion against rainy days and sustain future consumption. One way to test for precautionary saving is to estimate a consumption Euler equation, regressing consumption growth on the volatility of future consumption growth:¹

$$(1) \quad \Delta \ln c_{it} = \alpha V_{it}^c + \beta' X_{it} + \zeta_{it},$$

where c_{it} is individual i 's consumption, V_{it}^c consumption growth volatility conditional on information at the beginning of the period, X_{it}' a vector of demographic controls and ζ_{it} a consumption growth innovation. Under CRRA preferences, α equals 1/2 the coefficient of relative prudence as defined by Kimball (1990) and thus identifies the strength of the precautionary motive for saving.

I. Problems with and solutions to the Euler equation approach

Dynan (1992) was one of the first to use (1) to identify the strength of the precautionary motive using data from the Consumer Expenditure Survey. The main problem is that V_{it}^c is unobserved. Dynan solves it by estimating:

* Andreas Fagereng: Statistics Norway, Research Department, Postboks 8131 Dep., NO-0033 Oslo, Norway (email: afagereng@gmail.com); Luigi Guiso: Einaudi Institute for Economics and Finance, Via Sallustiana 62, 00187 Rome, Italy (email: guiso@tin.it); Luigi Pistaferri: Stanford University, Department of Economics, Stanford, CA 94305-6072, USA (email: pista@stanford.edu). We are grateful to The Research Council of Norway for financial support (grants #236935 and #230843).

¹An alternative approach is to estimate reduced form savings equations or wealth equations using measures of earnings volatility, as in Guiso et al. (1992).

$$(2) \quad \Delta \ln c_{it} = \alpha (\Delta \ln c_{it})^2 + \beta' X_{it} + \zeta_{it},$$

and instrumenting realized volatility in consumption growth (the term $(\Delta \ln c_{it})^2$) with a set of socio-demographic variables that are known at time $t - 1$ and presumably explain volatility (such as education, occupation, etc.). She estimates a very small value of α , suggesting a limited role for precautionary saving. The empirical problem with this approach is that the selected instruments either lack power or may be good predictors of changes in the chosen pattern of consumption (i.e., they violate the exclusion restriction).²

An alternative approach is to proxy the unobserved conditional variance of consumption growth (a sufficient statistic for all sources of risk faced) with income risk, which for most people is the most relevant risk faced over the life cycle. The strength of this proxy depends on the degree of self-insurance. Suppose that consumption risk V^c originates from two independent sources, the first related to labor income risk, V^y , and the second, V^η , reflecting other sources of consumption risk arising for instance from shocks to consumption needs:

$$(3) \quad V_{it}^c = \theta V_{it}^y + V_{it}^\eta$$

The parameter $0 \leq \theta \leq 1$ maps income risk into consumption risk. A value of $\theta < 1$ implies partial self-insurance of labor income risk, either through accumulation of precautionary as-

²One solution to the problem of weak instruments, first proposed by Bertola et al. (2005), is to instrument the realized variability of consumption growth with a measure of the subjective variance of future earnings available in the Italian Survey of Households Income and Wealth. This approach results in an estimate of the prudence parameter of around 2, suggesting a strong precautionary motive. While useful, the problem with this approach is that subjective probability distributions of future earnings are usually unavailable in survey data. Furthermore, little is known on how much of the earnings risk that people face is actually insured.

sets or other sources of insurance, such as family risk sharing or formal access to credit or insurance markets. We are interested in identifying both the strength of the precautionary motive α as well as the extent of insurance against labor income risk θ . The latter is critical to judge the adequacy of precautionary savings.

Using income risk as a proxy for the unobserved V_{it}^c leads to the estimation of a modified Euler equation:

$$(4) \quad \Delta \ln c_{it} = \alpha \widetilde{V}_{it}^y + \beta' X_{it} + u_{it},$$

where \widetilde{V}_{it}^y is a measure of the *observed* residual variance of earnings faced by individual i at time t (obtained from estimation of a statistical process for earnings in longitudinal data). Since from (3) consumption volatility reflects the underlying volatility of individual earnings, the latter is a natural candidate to proxy for the former; it is idiosyncratic and can be measured. However, because a non-negligible part of measured earnings volatility (as discussed in Fagereng et al., 2016) may reflect choice rather than risk, it may fail the exogeneity requirement: it is an error-ridden measure of true earnings risk. To see this, suppose that (residual) wage variance is:

$$(5) \quad \widetilde{V}_{it}^y = H_{it} + \rho_f F_{it}$$

where F is the variance of shocks to the firm the individual works for, ρ_f the (square) of the pass-through of firms shocks to workers' wages reflecting partial insurance within the firm, and H the residual, non-firm related, earnings variance (e.g., the unobserved component of human capital). However, assume that only a fraction ρ_h of H is truly risk, while the rest is variation in wages that reflects individual choice (or even statistical noise), rather than risk. Hence:

$$(6) \quad V_{it}^y = \rho_h H_{it} + \rho_f F_{it},$$

implying that measured earnings risk is $\widetilde{V}_{it}^y = V_{it}^y + \varepsilon_{it}$, the sum of true risk and a measurement error $\varepsilon_{it} \equiv (1 - \rho_h) H_{it}$. An OLS regression of Δc_{it} on \widetilde{V}_{it}^y as in (4) thus yields:

$$p \lim \widehat{\alpha} = \alpha \theta \frac{\rho_h V_H + \rho_f^2 V_F}{V_H + \rho_f^2 V_F}$$

where $V_H = \text{var}(H)$, $V_F = \text{var}(F)$ (see Online Appendix, OA). This is a downward biased estimate of α for two reasons: i) it ignores that some earnings volatility may reflect choice or noise (the factor ρ_h); and ii) it mixes precautionary motive and self-insurance (the factor θ).³

Identification of the strength of the precautionary motive and of the self-insurance of earnings risk requires separate identification of the two parameters α and θ ; it also requires solving the error-in-variable problem induced by the choice/noise issue. Our proposed solution is to follow Fagereng et al.'s (2016) methodology, match data on consumers with data on their employers, and use the firm-specific variance of value added shocks (F_{it} in the notation of equation 6) as an instrument for \widetilde{V}_{it}^y in equation (4), and as an instrument for realized consumption volatility in equation (2). The instrument is valid because, as we document below, firms pass some of their productivity shocks onto wages (i.e., there is only partial insurance within the firm). Moreover, for most workers firm-related shocks cannot be manipulated and are hard to avoid - that is they are a genuine source of risk.⁴ It can be shown that, under the maintained assumptions, the resulting IV estimation of (4) provides a consistent estimate of $\alpha\theta$, while IV estimation of (2) provides a consistent estimate of α (see OA). Hence, comparison of estimates from the two models allows separate identification of α and θ .

II. Data

We use nearly 20 years of longitudinal administrative data on income and assets covering the whole Norwegian population. Data are thus attrition-free except for mortality and migration. A full description of the data is provided in Fagereng et al. (2016). Here we summarize their main features. All data, including

³Banks et al. (2001) interact the residual earnings variance with a measure of θ (the squared wealth/income ratio), and hence their estimate is free from the second problem.

⁴Note that if only because of self-insurance, it is natural to expect firm volatility to be a more powerful instrument for residual earnings volatility than for realized consumption volatility.

assets values and income from capital and labor, are obtained from tax administrative records implying that measurement error is likely contained. Firm data also come from administrative records (balance sheets). Employed workers can be matched with their employer and with measures of the firm performance for all years they are observed in employment. Because we use firm performance volatility as an instrument for earnings volatility we limit our sample to individuals employed in private firms.

A. Measuring consumption

The Norwegian administrative data do not provide us with a direct measure of consumption. Instead, we follow Fagereng and Halvorsen (2015) and impute it using the budget identity: $c_{it} = (y_{it}^d - \Delta a_{it})/n_{it}$, where y_{it}^d is the sum of all income sources (including capital income) net of taxes, and a_{it} is liquid assets (since information on housing wealth is incomplete). We deflate consumption by the OECD equivalence scale $n = (1 + 0.7(A - 1) + 0.5K)$, where A is the number of adults and K the number of children (aged less than 18) in the household. This measure adjusts for capital gains and losses by use of broad domestic and foreign stock market indices, but likely suffers from some measurement error. Examples are extreme observations that may occur in household-year observations where the household has been involved in a real estate transaction, extreme returns from financial markets or when the household is a business owner or a farmer. Excluding such observations and using the crude adjustment for capital gains and losses alleviate the problem of measurement error. However, issues related to the yearly nature of our data, such as intra-year trading in stocks, and to housing transactions and marriage/divorce dynamics, remain. For this reason, we focus on a sample where we exclude extreme values of consumption growth.⁵

⁵We choose an asymmetric trimming, excluding the bottom 2.5% and top 1% of observations because our measure of consumption is total expenditures, including durable purchases. Hence, occasionally one sees big increases in expenditures. To further limit imputation error we retain only households with consumption levels above a minimum threshold (around USD 10,000 following the government guidelines for social financial support).

III. Methodology

We measure firm j performance with its value added, VA_{jt} , 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 stochastic 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 .

We model the earnings y_{ijt} (in logs) of worker i in firm j as a linear function of a vector of workers' observed characteristics, \mathbf{Z}_{ijt} , and an idiosyncratic unexplained residual ω_{ijt} . The latter is the sum of an individual random walk and a transitory component, plus a component that depends on the firm shocks (transitory and permanent) with transmission coefficients λ^T and λ^P , reflecting partial wage insurance within the firm.⁶ Hence:

$$\begin{aligned}\ln y_{ijt} &= \mathbf{Z}'_{ijt}\gamma + \omega_{ijt} \\ &= \mathbf{Z}'_{ijt}\gamma + v_{it} + \lambda^P f_{jt}^P + \lambda^T f_{jt}^T\end{aligned}$$

where v_{it} is an individual unobserved component (the sum of a random walk permanent component and an i.i.d. transitory shock). The term v_{it} may in part be predictable by the consumer (though not by the econometrician) and not constitute a source of earnings risk. The firm component instead qualifies as earnings risk. We assume the two components are independent. Hence, if earnings risk \tilde{V}_{it}^y is measured by the variance of ω_{ijt} it will reflect both true risk and "noise". For firm-related shocks to matter, λ^T and λ^P must be economically and statistically significant: i.e., firms must pass over to the workers some of the shocks to their performance and not offer them full wage insurance.

Following Guiso et al. (2005) methodology, Fagereng et al. (2016) use Norwegian data

⁶These processes fit the data quite well (see Fagereng et al. (2016)).

and obtain estimates of λ^P and λ^T of 0.071 and 0.018, respectively, both strongly statistically significant. Thus, both transitory and permanent shocks to firm performance are partly passed over to the worker but the pass-through is much stronger for permanent shocks, consistent with intuition. We replicate this methodology here.

To have a reasonably long series of earnings volatility measures, we compute the overall variance of unexplained workers earnings growth over T periods using rolling averages: $\widetilde{V}_{it}^y = \frac{\sum_{s=0}^{T-1} \omega_{ijt-s}^2}{T}$ (and set $T = 5$). We use this measure when estimating the Euler equation (4). Our instruments are the variances of the unexplained firm value added growth - both permanent and transitory - computed over the same T periods:

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

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

where $g_{jt} = \Delta(\ln V A_{jt} - \mathbf{X}'_{jt}\phi)$.

IV. Results

Table 1 shows the results of the IV estimates of the Euler equations (4) and (2).⁷ To eliminate outliers, we trim the variances of firm shocks at the 1st and 99th percentile. Our estimation sample contains about 327,000 observations. As additional controls we include a set of year fixed effects (to account for changes in interest rates), a quadratic in age, the change in the number of children and years of schooling.

The first column shows the estimates of the earnings-variance Euler equation (eq. (4)). The effect of earnings volatility is estimated to be 0.65, and is highly statistically significant. This is joint evidence of a precautionary savings motive and of consumers' inability to fully self-insure wage risk. The F test from the first stage (value 294.9) reveals that the instruments are

quite powerful. Table A2 (in the OA) shows the full first stage regression. The second column shows the estimate of the consumption-variance Euler equation (eq. (2)). The IV estimate is now 1.01, significant at the 1% level and larger than the estimate in the first column. As anticipated, the power of the instruments is lower when estimating this Euler equation than the one in the first column, though the test (value 7.8) is not far from conventional levels denoting non-weak instruments. In fact, in the first stage regression both instruments have predictive power on observed consumption volatility, but much more the variance of persistent shocks to the firm than that of the transitory ones (see OA). This is reasonable as the former are more easily self-insured than the latter.

From the estimates reported in the second column we infer a degree of prudence - measuring the strength of the precautionary motive - of around 2, which is reasonable. Using jointly the estimates reported in the first and second column, we infer a value of $\theta = 0.64$ - i.e. little less than $2/3$ of the earnings risk results in undesired fluctuations in consumption, while consumers manage to self-insure the remaining $1/3$.

Finally, we can use our estimates to assess the precautionary savings response to an increase in earnings risk when the latter originates from a change in the role of the firm as an insurance provider. Suppose that, following a trend documented by Benabou and Tirole (2016), firms offer more high-powered wage contracts and start sharing permanent shocks equally with their workers. That, is the value of λ^P increases from 0.07 to 0.5. Holding constant the self insurance parameter at 0.64, we calculate that consumption growth would be higher by 0.45 percentage points. If the firm shares equally also transitory shocks, the consumption profile would increase at a much faster rate of 0.8 percentage points.

V. Conclusion

Building on a credible instrument for consumption risk, we develop a strategy that allows to identify simultaneously the strength of the precautionary motive and the degree of self-insurance of labor income risk. At the same time, it provides a framework for studying the precautionary savings response of structural changes in wage insurance provided by the firm.

⁷Table A1 (in the OA) provides summary statistics of the main variables in the analyses. It is worth noticing that the variance of consumption growth and the variance of residual income are not comparable. While the latter measures variability in income after filtering the predictable component, the former reflects both predictable and unpredictable variation.

TABLE 1—IV ESTIMATES OF THE EFFECT OF WAGE AND CONSUMPTION RISK ON THE GROWTH OF CONSUMPTION

Endogenous:	$\Delta \ln c_{it}$	
	\tilde{V}_{yit}	\tilde{V}_{cit}
IV:	0.650*** (0.233)	1.008*** (0.310)
F-stat 1st stage	294.854	7.853
Kleibergen-Paap underid p-value	0.000	0.001
Hansen J-test p-value	0.029	0.396
Observations	327,518	327,518

Notes: The table reports IV estimates (see Table A2 in the OA for full first stage regressions) of the marginal effect of wage and consumption risk on the growth of consumption, using two instruments - the variance of transitory and permanent shocks to firm's value added. All regressions include year fixed effects, a quadratic in age, the change in the number of children, and years of schooling. F-stat for the power of the instruments is 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 a strong precautionary motive, a partial ability to self-insure labor income risk and a large precautionary savings response to firm adoption of high powered wage contracts.

REFERENCES

- Banks, James, Richard Blundell and Agar Brugiavini, (2001) "Risk Pooling, Precautionary Saving and Consumption Growth", *Review of Economic Studies*, 68, 4, 757-779.
- Benabou, Roland and Jean Tirole (2016), "Bonus Culture: Competitive Pay, Screening, and Multitasking", *Journal of Political Economy*, 124, 2, 305-370.
- Bertola, Giuseppe, Luigi Guiso and Luigi Pistaferri (2005), "Uncertainty and consumer durables adjustment", *Review of Economic Studies*, 72, 973-1007.
- Carroll, Christopher D. (1992), "The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence", *Brookings Papers on Economic Activity*, Volume 1992, Issue 2, 61-156.
- Dynan, Karen (1992), "How Prudent are Consumers?", *Journal of Political Economy*, 101, 6, 1104-1113.
- Fagereng, Andreas, Luigi Guiso and Luigi Pistaferri (2016), "Portfolio Choices, Firm Shocks and Uninsurable Wage Risk", NBER Working paper No. 22883.
- Fagereng, Andreas and Elin Halvorsen (2015), "Imputing consumption from Norwegian income and wealth registry data", *Journal of Economic and Social Measurement*, forthcoming.
- Guiso, Luigi, Tullio Jappelli and Daniele Terlizzese (1992), "Earnings uncertainty and precautionary saving", *Journal of Monetary Economics*, 30, 2, 307-337.
- Guiso, Luigi, Luigi Pistaferri and Fabiano Schivardi (2005), "Insurance within the Firm", *Journal of Political Economy*, 113, 5, 1054-1087.
- Kimball, Miles (1990), "Precautionary Saving in the Small and in the Large," *Econometrica*, January 1990, 58 (1), 53-73.