Crawford School of Public Policy

TTPI

Tax and Transfer Policy Institute

Risk aversion among Australian households

TTPI - Working Paper 2/2019
February 2019

Robert Breunig
Crawford School of Public Policy, Australian National University

Owen Freestone
Crawford School of Public Policy, Australian National University

Abstract

This paper explores risk aversion among Australian households using panel data from the Household Income and Labour Dynamics in Australia (HILDA) survey. Using households' share of risky assets, we test whether relative risk aversion is constant in wealth. After accounting for measurement error, we cannot reject the constant relative risk aversion (CRRA) assumption. Using an Euler equation that adjusts for measurement error in consumption data, we estimate the coefficient of relative risk aversion in the CRRA utility function. Point estimates from our preferred non-linear models suggest a moderate degree of risk aversion for the typical Australian household, with values ranging from 1.2 to 1.4. These findings can provide guidance for calibrating household preferences in macroeconomic models of the Australian economy.

JEL Codes: D12, D15, E21
Keywords: measurement error, intertemporal consumption choice, Euler equations, risk aversion, GMM, instrumental variables

∗*Author contact details: owen.freestone@anu.edu.au; robert.breunig@anu.edu.au*

Tax and Transfer Policy Institute

Crawford School of Public Policy

College of Asia and the Pacific

+61 2 6125 9318

tax.policy@anu.edu.au

The Australian National University

Canberra ACT 0200 Australia

www.anu.edu.au

The Tax and Transfer Policy Institute in the Crawford School of Public Policy has been established to carry out research on tax and transfer policy, law and implementation for public benefit in Australia. The research of TTPI focuses on key themes of economic prosperity, social equity and system resilience. Responding to the need to adapt Australia's tax and transfer system to meet contemporary challenges, TTPI delivers policy-relevant research and seeks to inform public knowledge and debate on tax and transfers in Australia, the region and the world. TTPI is committed to working with governments, other academic scholars and institutions, business and the community.

The Crawford School of Public Policy is the Australian National University's public policy school, serving and influencing Australia, Asia and the Pacific through advanced policy research, graduate and executive education, and policy impact.

1 Introduction

Risk aversion is central to the theory of individual decision-making under uncertainty underpinning models of inter-temporal consumption and saving behaviour, portfolio choice and labour supply to name a few. Risk aversion has also been identified as an important empirical factor in a broad range of economic and social choices.[1] The household's degree of risk aversion is also a crucial input in large-scale macroeconomic models, where it plays a large part in determining the economy's steady state and the dynamic response of households to changes in interest rates, taxes, and other aspects of the economic environment.[2]

While a range of international studies, mostly from the U.S., have estimated the key parameter which determines risk aversion, there is no evidence, to our knowledge, for Australia. As it seems plausible that Australian households might differ substantially from their U.S. counterparts in their attitude towards risk and therefore their inter-temporal consumption and portfolio choices, it is important to estimate this parameter for Australia. This will be particularly useful for researchers who build large-scale models of the Australian economy. This is our main contribution in this paper.

We provide Euler equation-based estimates of the risk aversion parameter (γ) from the constant relative risk aversion (CRRA) utility function for Australia using longitudinal consumption data from the Household Income

[1]While too numerous to list here, an interesting example is an individual's occupational choice, where more risk averse individuals have been found to choose careers in occupations that involve lower earnings risk (Bonin et al., 2007).

[2]The related concept of 'prudence' is also an important factor in decision-making under uncertainty. See Kimball (1990) on this point.

and Labour Dynamics in Australia (HILDA) survey and data on aggregate interest rates. These are the first Euler equation estimates using micro data in Australia as far as we are aware.

To preview the results, our preferred non-linear specifications that allow for the well-known and substantial problem of measurement error suggest γ is in the range of 1.2 to 1.4. This implies a relatively modest degree of risk aversion for the mean household, similar to the 1.5 found in Alan et al. (2009), a recent comparable US study.

The CRRA utility function and our estimation approach embed the assumption that relative risk aversion is constant with respect to a household's level of wealth. Previous Australian studies have found mixed evidence regarding this assumption. We also test this assumption and we are unable to reject the assumption of constant relative risk aversion. We also attempt to explain the variety of findings from previous Australian studies.

In the next section, we provide some background about the research agenda around estimating risk aversion. Section 3 describes the data we use while section 4 details our tests of the assumption that relative risk aversion is constant in wealth. Having established some empirical support for the CRRA assumption, section 5 presents our γ (and other parameter) estimates from various Euler equation models derived from the CRRA assumption. Section 6 provides some conclusions and areas where further research would be useful in the Australian context.

2 Background

Given the fundamental interest around risk aversion and its important in economic modeling, an extensive body of empirical research attempts to quantify the degree of risk aversion among households. Although there are a number of possible ways to proceed,[3] this paper follows the 'Euler equation' approach pioneered in Hall (1978). In basic terms, this approach exploits the relationship between consumption and interest rates derived from the household's inter-temporal optimisation problem to uncover the risk aversion parameter in the household's underlying utility function.

To operationalise the Euler equation approach, a particular form of preferences must be assumed, and we opt for the constant relative risk aversion (CRRA) utility function $U(C) = \frac{C}{1-\gamma}^{1-\gamma}$, where C is consumption and γ is the coefficient of relative risk aversion. A particular advantage of CRRA preferences is that they imply an empirically tractable Euler equation, namely:

$$\mathrm{E}_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} (1 + r_{t+1})\beta \right] = 1 \qquad (1)$$

where r_t is the return on a generic asset at time t, β is the household's discount factor and the expectation operator E_t indicates that the household's consumption decision takes place in an environment where the value of future variables (such as earnings) is uncertain.

[3] A more structural approach would be to estimate the degree of risk aversion via simulated method of moments, as done in Gourinchas and Parker (2002). Alternatively, stock returns can be used to infer the degree risk aversion, which is the approach underlying the 'equity premium puzzle' literature (see Kocherlakota (1996) for example). A completely different approach to estimating individual risk aversion is employed in experimental studies such as Levy (1994).

CRRA preferences have a number of other desirable properties, which have made them a popular specification in models of inter-temporal consumption decisions under uncertainty. This includes the convenient feature that the standard (Arrow-Pratt) measure of relative risk aversion is constant in wealth and given directly by γ,[4] while the elasticity of inter-temporal substitution (EIS) is also constant and equal to $1/\gamma$.[5] CRRA preferences also have a number of economically appealing aspects, including the presence of 'prudence', which gives rise to a precautionary (or 'buffer stock') savings motive when households face future income uncertainty.[6]

The Euler equation approach that we use was initially discredited in the early 2000s in the United States literature because it generally failed to deliver consistent and reliable estimates of γ and β despite the large number of studies (Carroll, 2001a). This disappointing performance was attributed to a number of factors including measurement error in consumption data and difficulty in finding suitable instruments for the endogenous variables, especially in the case of the log-linearised Euler equation as explained below.

More recently, lengthening panel datasets, improved availability of consump-

[4]For a generic utility function $U(C)$, the Arrow-Pratt measure of relative risk aversion is given by -$CU''(C)/U'(C)$. For other preference specifications, the degree of relative risk aversion and the elasticity of inter-temporal substitution (EIS) are not simple mappings from preference parameters, and instead depend on the current values of consumption (or wealth).

[5]This simple inverse relationship for CRRA preferences is what allows the Euler equation to identify the agent's degree of risk aversion. However, there are more general forms of preferences in which risk aversion and the EIS are governed by separate parameters. Encouragingly, the one study to test this inverse relationship empirically found in favour of it (Yagihashi and Du, 2015).

[6]For CRRA preferences, relative prudence is equal to γ+1. In general, a utility function exhibits prudence when its third derivative with respect to consumption is positive (Kimball, 1990).

tion data, and analytical refinements have allowed estimation of γ and β, using the Euler equation aproach, in a more precise way. For example, Alan et al. (2009) derive a GMM panel estimator that explicitly allows for measurement error in consumption data, while Alan et al. (2018) show that synthetic panel data based on repeated cross-sections can overcome the endogeneity problems given a long enough time dimension.[7]

Following these recent refinements in the international literature, we provide Euler equation-based estimates of γ for Australia using data from the Household Income and Labour Dynamics in Australia (HILDA) survey. Based upon empirical research and simulation studies from other countries, the HILDA data is only now becoming long enough to support this kind of analysis, with around 14 years considered to be the minimum required panel length (Alan et al., 2009).

While CRRA preferences are a convenient assumption, they impose a restriction on the relationship between risk aversion and wealth, namely that relative risk aversion is constant with respect to a household's level of wealth. In fact, the empirical support for this form of preferences is not overwhelming, in part because there are relatively few studies that have formally tested the CRRA assumption.[8] Therefore, we undertake an initial test of whether the CRRA assumption is in fact a good one for Australian households, using an approach based on a household's share of risky assets, similar to that in

[7]Both studies use evidence from Monte Carlo simulations to demonstrate the reliability of their proposed estimators using plausible data generating processes and based on realistic features of the data with regards to panel length and the extent of measurement error in consumption.

[8]Some examples include Ogaki and Zhang (2001); Chiappori and Paiella (2011); Tsigos and Daly (2016); Conlin et al. (2016).

Chiappori and Paiella (2011) and Tsigos and Daly (2016) (the latter being a recent Australian study). Once we account for endogeneity in risk aversion and wealth as well as measurement error – which we show is crucial in such regressions – we cannot reject the CRRA assumption.[9] While constant for a particular household over time, our analysis reveals considerable variation in relative risk aversion *between* households with a median γ of 1.8.[10] As a cross-check, we find that the distribution of this 'objective' measure of risk aversion is significantly correlated with the two subjective measures of risk aversion contained in the HILDA survey.

We next turn to a detailed description of our data and sample selection.

3 Data

The HILDA survey began in 2001 and provides annual data on a broad range of economic and social topics. The samples used in all components of our analysis comprise household heads of working age. A head is defined as the oldest male member of a household, or the oldest female in households without a male adult.

Our definition of household head is not important for the results presented below. We identify a 'household head' so that we can follow individuals

[9]In Appendix C, we show that failing to take account of measurement error in wealth in a regression of risky asset share on wealth can result in a negative estimated coefficient on wealth even when the true coefficient is known to be zero.

[10]While the median value of 1.8 is somewhat above our Euler equation estimates, we note below estimates of γ using risky asset shares are very sensitive to the definition of risky asset (see Appendix B) and assumptions about the market portfolio. For this reason, we consider the point estimates from our Euler equation analysis to be a better indicator of γ for the 'average' or 'typical' household.

over time using the `xwaveid` variable in HILDA. The analysis is based upon household-level data so choice of head is rather arbitrary. In our Euler equation estimates below we only use couples that stay together so choosing the female as the head of household does not alter the results. We have verified this using alternative definitions of household head and the results presented below are unchanged.

We use the sum of expenditure on groceries and meals out as a proxy for household consumption. Unlike other components of household expenditure, information on these components has been collected in almost all years of the survey.[11] Just as importantly, the data quality of these items is probably higher than other components of expenditure in the HILDA survey (Wilkins and Sun, 2010), although they still contain significant measurement error as we quantify in Appendix D. These nominal series are deflated using State-specific consumer price indices for food.

Because of differing data requirements, the samples used in our risky asset share and Euler equation estimates differ in a number of key respects, as we now describe.[12]

Risky asset share regressions

The risky asset share regressions rely on HILDA's wealth module which is included in only four survey years: 2002, 2006, 2010 and 2014. The sample is based on an unbalanced panel of these four years consisting of all responding

[11] The expenditure information is taken from the household questionnaire for the years 2001, and 2003 to 2005, and from the self completion questionnaire for subsequent years.

[12] Further details are provided in Appendix A.

household heads between ages 23 and 59 inclusive.

In the base model, risky assets are defined narrowly as equities, to allow comparability with the two exiting Australian studies, although we consider broader definitions of risky assets in extended models. The regressions only include households with positive risky assets (however defined), and those with information on household size and age of household head. The mean value of the risky asset share is 17 per cent based on the narrow definition of risky assets. The instrumental variable (IV) regressions also require non-missing data on the instruments, that is, household disposable income and household consumption (which we again proxy using expenditure on groceries and meals out).

Consumption Euler equations

Again, our Euler equation analysis uses grocery and meals out expenditure as a proxy for overall household consumption, an approach also adopted in most overseas studies.[13] The panel encompasses the period 2001 to 2016 inclusive, except for 2002 where information on expenditure was not collected in HILDA.

Our sample selection choices closely follow Alan et al. (2009) to aide comparability with that study. We restrict the sample to households consisting of couples who were in a stable relationship over the period in which they feature in the survey. Also, because the Euler equation only holds for an in-

[13] This includes the recent US study by Alan et al. (2009). The use of food (and meals out) expenditure as a proxy for consumption in most US studies is because food was the only component of household expenditure covered in the US Panel Survey of Income Dynamics (PSID) from the survey's inception in 1968 until 1999.

terior solution, we exclude households who are liquidity constrained, which we define as having zero financial assets in any of the years they feature in the panel. The latter restriction removes only 36 observations. The real interest rate series is calculated as the annual-average nominal 3-month bank accepted bill rate less inflation expectations pertaining to that period.[14]

4 Testing the CRRA assumption using share of risky assets

This section reviews the empirical support for CRRA preferences before we proceed to our Euler equation-based estimates. We closely follow Chiappori and Paiella (2011), a recent study that analysed whether relative risk aversion is constant using Italian panel data on risky asset shares. By solving a 2-period optimisation problem of a risk-averse investor, they show that the share of wealth invested in risky assets is proportional to investor h's relative risk aversion, γ_h:

$$\alpha^h = \frac{1}{\gamma^h} \frac{\mathrm{E}[r_m - r_f]}{\sigma_m^2} \tag{2}$$

where α_h is the risky asset share, r_m is the return on the risky asset, r_f is the risk free rate and σ_m^2 is the variance of the return on the risky asset.

However, they also formally demonstrate that cross-sectional wealth data alone is insufficient to test whether relative risk aversion is constant. This

[14]We experiment with two different measures of inflation expectations: the Reserve Bank of Australia's survey of market economists and the inflation rate implied by inflation-indexed bond prices. The key results are not greatly affected by this choice.

is because the distribution of risky asset share in the population depends on the joint distribution of wealth and risk aversion, not just the form of preferences (Chiappori and Paiella, 2011). For example, if risk aversion is heterogeneous across the population, it is reasonable to expect that less risk averse people will earn, on average, higher returns and accumulate greater wealth over time. This would lead to a negative correlation between relative risk aversion (measured by risky asset share) and wealth cross-sectionally, even when the underlying preferences were CRRA. Chiappori and Paiella show that panel data can overcome this problem, by allowing one to test whether relative risk aversion varies with wealth for a particular household.

In Australia, Tsigos and Daly (2016) test the relative risk aversion assumption using HILDA wealth data. Using a different measure of relative risk aversion (based on observed portfolio weights for each household), Tsigos and Daly find a negative correlation between risk aversion and wealth. However, they find no relationship between relative risk aversion and wealth when they instead use risky asset share as their measure of risk as in Chiappori and Paiella (2011). In contrast, Cardak and Wilkins (2009) find a small positive relationship between risky asset share and wealth. Their study was also based on the HILDA survey, but they only had access to a single cross-section of wealth data (for 2002) that was available at the time.

Given the mixed findings in Australia, we re-test the question of whether relative risk aversion varies with wealth. Compared to the previous Australian studies by Cardak and Wilkins and Tsigos and Daly, we benefit from an additional 3 waves and 1 wave of data respectively. In particular, this allows us to use the panel dimension of the HILDA survey, which was unavailable to Cardak and Wilkins. We follow the measure of risk aversion used in Chiappori and Paiella and Cardak and Wilkins based on risky asset shares.

The basic equation that we estimate is:

$$\log(\alpha_t^h) = \beta_0 + \beta_1 \log(W_t^h) + \beta_2 X_t^h + u^h + \nu_t^h \tag{3}$$

where α_t^h is the share of risky assets for household h at time t, W_t^h is total financial wealth, u^h is a fixed effect and ν_t^h is random error term. To control for the likely endogeneity between risky asset share and wealth, we also estimate the equation in first differences (FD) which eliminates the possible bias from u^h:

$$\Delta \log(\alpha_t^h) = \beta_0 + \beta_1 \Delta \log(W_t^h) + \beta_2 \Delta X_t^h + \Delta \nu_t^h \tag{4}$$

As pointed out in Chiappori and Paiella, measurement error in the wealth data can also result in biased estimates of the risk aversion-wealth relationship if not adequately dealt with. This is because total financial wealth is the denominator of the dependent variable and therefore appears on both sides of equation 4. We show in Appendix C that for a mean risky asset share of less than 50 per cent, standard multiplicative measurement error in the wealth data will result in moderate to severe downward bias in estimates of β_1 given plausible values for the variance of the measurement error.[15]

[15]In Appendix C, we assume that 20 per cent of the variance in observed wealth is measurement error. While there is very little evidence on the extent of measurement error

To address the problem of measurement error, we instrument wealth with disposable household income and food consumption in the levels equation. Similarly, we use growth in disposable income and growth in consumption as instruments for wealth in the first difference equation. Our final specification also includes year dummies (to control for macroeconomic effects) and controls for household composition and age.

4.1 Results

As suspected, the estimated relationship between risky asset share and wealth is sensitive both to whether we use instrumental variables to adjust for the impact of measurement error and whether the model is estimated in levels or first differences (Table 1). The basic OLS regression in levels (column 1) shows a strong and highly significant negative relationship (-0.13) between risky asset share and wealth. However, once we adjust for the impact of measurement error by instrumenting for wealth, the coefficient on wealth falls in magnitude and becomes insignificant (column 2). As we argued above, there are good reasons to think that measurement error would contribute to a negative cross-sectional relationship even when none truly exists when the mean risky asset share is less than 50 per cent.[16]

in wealth survey data, studies that have examined measurement error in food consumption suggest that this assumption is conservative (see for example Ahmed et al. (2006); Brzozowski et al. (2017)). Further, the bias caused by measurement error is likely to be exacerbated in the first difference regression, because the process of differencing data magnifies the noise-to-signal ratio.

[16]The mean share of risky assets in the data is well below 50 per cent unless superannuation is included. Including superannuation as a risky asset results in a mean risky asset share of 73 per cent. If the model in column 1 is re-estimated using this broader measure of risky assets, the estimated coefficient becomes close to zero, which is what we would expect based on our simulation graph in Appendix C.

Column 3 re-estimates the basic model in first differences. The coefficient on wealth is again negative, but is very imprecisely estimated and not significantly different from zero at standard confidence levels.

Columns 4 and 5 experiment with expanded definitions of risky assets. Column 4 adds cash investments (bonds and the like) and trusts to the measure of risky assets, while the regression in column 5 additionally includes business equity and superannuation. In the latter regression, the denominator is also adjusted to include business equity and superannuation. The coefficient on wealth remains insignificant, although it continues to be negative.

Given the difference between the OLS and IV estimates, we explore whether the results are robust to alternative instruments for wealth. One obvious alternative instrument for change in wealth is the cumulative difference between disposable income and consumption in the intervening periods between observations on wealth. In other words, this utilises income and consumption data in years that did not contain a wealth module in HILDA. For example, for 2006 we subtract cumulative food consumption from cumulative disposable income for each household over the period 2002 to 2006, which yields a 'flow' measure of change in wealth. When we use this variable as the instrument in models 3 to 5, the key results are unaffected.

We also conduct a weak instruments test, and find that our chosen instruments are well correlated with financial wealth and change in financial wealth.[17] A test of over-identifying restrictions fails to reject the null hy-

[17] The correlation between financial assets and income is 0.44 in levels and 0.18 in growth terms, while the correaltion between wealth and consumption is 0.26 in levels and 0.08 in growth terms.

Table 1: Risky asset share regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	OLS Level	IV Level	OLS FD	IV FD	IV FD	IV FD
Risky assets: numerator	E	E	E/FA	E	E+C+T	E+C+T+BE+S
denominator:	FA	FA	FA	FA	FA	FA+BE+S
Total Financial Assets*	-0.128	-0.047	-0.313	-0.173	-0.226	-0.033
	(0.0122)	(0.0452)	(0.0191)	(0.139)	(0.1355)	(0.0222)
Year dummies						
2006	-0.056	-0.08				
	(0.0405)	(0.042)				
2010	-0.229	-0.248	-0.157	-0.136	-0.099	-0.012
	(0.0437)	(0.0451)	(0.0642)	(0.0664)	(0.0658)	(0.0138)
2014	-0.363	-0.391	-0.168	-0.161	-0.166	-0.024
	(0.0463)	(0.0481)	(0.0677)	(0.069)	(0.0684)	(0.0135)
HH size	-0.205	-0.218	-0.258	-0.342	-0.211	0.035
	(0.0409)	(0.0418)	(0.0759)	(0.1018)	(0.0965)	(0.0195)
Age	0.076	0.077	0.015	0.012	0.008	0.007
	(0.0186)	(0.0189)	(0.0303)	(0.0319)	(0.0312)	(0.0056)
Age squared	-0.081	-0.085	-0.025	-0.022	-0.018	-0.01
	(0.0215)	(0.0219)	(0.0371)	(0.0388)	(0.0381)	(0.0071)
Constant	-1.415	-2.199	-0.16	-0.127	0.029	-0.055
	(0.3946)	(0.5738)	(0.6024)	(0.6331)	(0.6161)	(0.1064)
N	6952	6849	3111	3070	3245	10408

Standard errors in parentheses.
E=Equities; C= Cash investments; T=Trusts; BE=Business equity; S=Superannuation;
FA=Financial assets = E+C+T+ Bank accounts.
$*$ Equal to denominator used to calculate risky asset share.

pothesis that the chosen instruments are exogenous for models 3 to 5, but does reject the null for model 2. We also try including labour force participation as a component of Z_t, and use its lag as an instrument. This addition does not affect the main results in Table 1.

Comparison with Tsigos and Daly's results

In finding no significant relationship between risky asset share and wealth, our results confirm the results from Tsigos and Daly's sensitivity analysis but contradict their main result. However, their main result – that relative risk aversion falls with wealth – is based on a different measure of risk that uses portfolio weights calculated using external information on asset returns. Tsigos and Daly do not attempt to explain why the two approaches yield different results in their paper. It is also worth noting that, compared with Tsigos and Daly, our test of the risk aversion-wealth relationship benefits from an extra year of wealth data, giving our test additional power.

While we are unable to reject the CRRA assumption, the estimated coefficient on wealth is negative in all five specifications albeit insignificant in four cases. We therefore view our results as providing tentative support for the CRRA, particularly given that the insignificant estimates only arise in the instrumental variable regressions. Additional years of HILDA data going forward will help us to work out whether risk aversion really is constant with respect to wealth or not.

Distribution of implied relative risk aversion parameter

Equation 2 can be used to calculate the distribution of γ_h across the population. However, as Chiappori and Paiella (2011) note, this distribution is only identified up to a scale factor, given by the ratio of the excess return to the variance of the risky portfolio $\frac{r_m - r_f}{\sigma_m^2}$. Taking the plausible (and simple) case where this ratio is equal to one, the value of γ_h is just the inverse of the risky asset ratio for investor h. Clearly, the derived value of γ_h will depend inversely on the definition of 'risky asset', and will be higher or lower for narrower and broader definitions respectively. Additionally, very small values of α_h will result in very large values of γ_h. Following Chiappori and Paiella we therefore truncate the distribution of risky asset shares to exclude values less than 6 per cent and compute the median rather than mean value. The median estimates of γ_h are 1.1, 2.2 and 2.4 across the three definitions of risky assets used in our regressions. We also find little evidence that γ_h varies by wealth. Table B.1 in Appendix B contains further details.

Correlation with subjective data on risk aversion

An interesting question is whether the risk aversion measure derived from risky asset shares correlates with subjective measures of risk aversion contained in the HILDA survey. In particular, the two seperate survey questions ask respondents to indicate the degree of financial risk they are prepared to take.[18] The risky asset share is significantly correlated with both subjective

[18]The HILDA codes are `_firisk` and `_pntrisk` respectively. For the former variable higher values indicate higher risk aversion, while for the latter, higher values imply lower risk aversion.

measures, with a correlation coefficient of -0.25 for one subjective measure (which measures risk aversion) and 0.05 for the second measure (which measures risk appetite).

5 Estimates of risk aversion using a consumption Euler equation

Based upon our tentative conclusion that relative risk aversion is constant with respect to wealth, we now explore Euler equation estimation based on CRRA preferences. The Euler equation approach of recovering preference parameters began with Hall (1978), and we follow the essence of this approach. In particular, with CRRA preferences, a time discount factor β and an interest rate r, the Euler equation linking consumption in period t to consumption in period *t+1* for household h is:

$$\left(\frac{C_{t+1}^h}{C_t^h}\right)^{-\gamma}(1+r_{t+1})\beta\exp\left(\theta\Delta Z_{t+1}^h+\Delta v_{t+1}^h\right)-1=\epsilon_{t+1} \tag{5}$$

$$\mathrm{E}_t\left[\epsilon_{t+1}\right]=0 \tag{6}$$

Where the expectation operator relates to the time dimension, implying that the mean of the expectational error, ϵ_{t+1} is equal to 0 for a given household over time rather than for the cross section.

Two other points about Equation 5 are worth making. First, this condition holds only for an interior solution for consumption, and will generally not hold for households subject to binding borrowing constraints. Second, Equation (5) is derived on the assumption that consumption and leisure are additively separable in the household's utility function. As a sensitivity test of our main results, we also consider the case where consumption and leisure are non-separable.

Exact non-linear equation versus the linearised version

Previous studies have typically proceeded in one of two ways: estimating the exact, non-linear Euler equation as it appears in 5 using generalized method of moments (GMM); or taking a first- or higher-order linear approximation of the Euler equation and using OLS or linear IV estimation. Both approaches have drawbacks.

To begin with, both approaches require instruments for the model's endogenous variables – consumption growth and the interest rate. In the non-linear case, economic theory implies that any variables in the consumer's information set at time t are valid instruments, including lags of consumption growth and the interest rate. However, Carroll (2001a) and Alan et al. (2009)) show that estimates based on the exact non-linear equation are biased when the consumption data contains measurement error, which is almost certainly the case in practice. These studies suggest that estimates of β will be particularly biased in the presence of measurement error, and this bias does not improve with a longer panel length (Alan et al., 2009).

The linearised approach overcomes the problem with measurement error[19], but creates its own problems. To see this, first note that the first-order

[19] With multiplicative measurement error, log-linearising renders the measurement error additive meaning that measurement error would no longer affect the consistency of parameter estimates.

approximation is given by:

$$\Delta \log(C_{t+1}) = \frac{1}{\gamma}[\log(\beta) + \log(1 + r_{t+1})\theta\Delta Z_{t+1} \tag{7}$$
$$+\Delta v_{t+1} + \log(1 + \epsilon_{t+1}) + k_t]$$

where k_t includes higher order terms remaining from the approximation process.

The problem is that the composite error term is now likely to be correlated with consumption growth and interest rates via the higher-order terms in k_t. Carroll (2001a) shows that these high-order terms are inherently endogenous in a standard model of life-cycle consumption behaviour with wage uncertainty, and that estimates of γ will be severely biased. A further disadvantage of using equation 7 is that β can no longer be recovered as it is subsumed into the constant along with the mean values of the measurement error and approximation errors.

Various solutions have been proposed to overcome the difficulties with each of these approaches. In the non-linear case, Ventura (1994), Chioda (2004) and Alan et al. (2009) show that explicitly taking account of measurement error in consumption can yield relatively accurate estimates of the preference parameters γ and β even in panels of only moderate length (that is, around 15 years of data). Each of these studies essentially proceeds by assuming that true consumption ($C_t^\star$) is subject to log-normal iid measurement error η_t with equal variance σ_η^2 across households, such that $C_t = C_t^\star \eta_t$. The modified Euler equation is:

$$\mathrm{E}_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}(1+r_{t+1})\beta\exp\left(\theta\Delta Z_{t+1}\right)-\exp(\gamma^2\sigma_\eta^2)\right]=0 \tag{8}$$

While this single moment equation does not allow β and σ_η^2 to be separately identified, Alan et al. (2009) show that an analogous moment condition for two-period apart consumption (C_{t+2}/C_t) can be used to identify σ_η^2 (and therefore β).

Alternatively, Alan et al. (2018) show that relatively precise estimates of γ can be obtained using the linearised model if a long enough time series on consumption is available. However, their paper suggests that 'long enough' equates to perhaps 40 years of data, a requirement that few if any panel surveys satisfy. As a practical alternative, they show that a synthetic panel formed from repeated cross-sectional surveys over a long enough period can deliver reliable estimates of γ using the linearised Euler equation.

Our approach

With (at most) 15 years of available consumption data in the HILDA survey, we opt for GMM on the exact non-linear Euler equation, allowing for measurement error in our preferred specifications. For reference, we also obtain an estimate of γ from a linearised version of the model, noting the caveats above in relation to biased estimates in small panels such as ours.

In dealing with measurement error, we consider two approaches. As we use a single moment equation, neither approach allows us to separately identify the discount rate β and the variance of the measurement error σ_η^2

without further moment conditions or an 'external' estimate of σ_η^2.[20] In the first approach, we estimate a single parameter, which is some unknown combination of β and σ_η^2. In the second approach, we use separate estimates of the variance of the measurement error σ_η^2, which enables us to recover an estimate of β. This analysis is reported in Appendix D.

The vector ΔZ_t can in principal consists of any variables that may affect the marginal utility of consumption, including endogenous variables such as labour supply choices (Attanasio and Low, 2004). While we experiment with other variables, our final specification for ΔZ_t only includes change in household size as in Alan et al. (2009). Changes in household size should have a large and direct effect on marginal utility, and this has been confirmed empirically (Attanasio and Low, 2004). We also include a 2008 year dummy in some specifications. This dummy would capture a possible structural break in the relationship between consumption growth and interest rates that may have occurred with the onset of the global financial crisis.

Empirical model

Taking all of the above gives the following moment conditions for each period t (where the household superscript has been suppressed) whose sample counterparts are used in GMM estimation:

[20] We also consider adding a second two-period apart moment condition, which would allow β and σ_η^2 to be disentangled (Alan et al., 2009). However, we found that the addition of two-period apart moments greatly reduced the precision of the parameter estimates, resulted in implausible values for γ and often led to non-convergence in our GMM estimation routine.

$$\mathrm{E}_t\left\{\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}(1+r_{t+1})\beta\exp\left(\theta\Delta Z_{t+1}\right)-\exp(\gamma^2\sigma_\eta^2)\right]X_t\right\}=0 \tag{9}$$

where X_t is the set of instruments at time t, as described in the next section.

As noted above, we also estimate a standard linearised Euler equation:

$$\Delta\log(C_{t+1})=\alpha+\frac{1}{\gamma}\log(1+r_{t+1})+\frac{\theta}{\gamma}\Delta Z_{t+1}+e_{t+1} \tag{10}$$

where α now includes $\log\beta$ as well as the means of the higher-order approximation errors, and e_t now includes the household's expectational error, the measurement error and the time-varying components of the approximation error.

Choice of instruments

We follow many previous studies in using lags of the interest rate (in addition to a constant) and consumption growth as instruments in our non-linear GMM estimation. There are a number of reasons why these are likely to be good instruments. First, rational expectations suggests that lagged variables, which are known to the household at time t, will be uncorrelated with their forecast error the following period. Second, consumption and interest rates tend to be adequately correlated with lags of themselves, avoiding the problem of 'weak instruments'. Lastly, simulation exercises have shown that they are valid in a standard life-cycle model environment (Alan et al., 2009). We also use change in household size as an instrument for itself, and do likewise with the 2008 year dummy where present. We also experiment with dropping lagged consumption growth from our instrument list (see column 4 in Table D.1). This leaves the lagged interest rate and a constant to identify the two parameters of interest γ and β, which is essentially the just-identified model applied in Alan et al. (2009). Because the linearised model is a log transformation of the true model, we substitute log versions of the instruments used in the non-linear equation.[21]

All estimates of the non-linear model are based on a standard two-step GMM with robust weight matrix. Estimation based on an alternative weight matrix, such as iterative GMM, had very little impact on the results. The linearised model is also estimated using (linear) GMM.

Allowing for non-separable consumption and leisure

The Euler equations (5) and (10) were derived on the assumption of separability between consumption and leisure. However, the few international studies that have formally tested separability have generally found against this assumption.[22] As a sensitivity test, we estimate a modified Euler equation derived from non-separable Cobb-Douglass preferences $U = \frac{1}{\rho_1\rho_2}(C^{\rho_1}L^{1-\rho_1})^{\rho_2}$, where L is leisure. Here, the coefficient of relative risk aversion is a combination of both utility parameters and is given by $1 - \rho_1\rho_2$. The resulting Euler equation is:

[21] In fact, by subsuming the measurement error into the error term, the error term becomes an MA(1) process. While this means that the first lag of the interest rate is no longer a valid instrument, we find that using the second lag of the interest rate results in implausible (negative) estimates of γ.

[22] We are not aware of any Australian studies that have examined this question.

$$\mathrm{E}_t \left[\left(\frac{C_{t+1}^h}{C_t^h} \right)^{\rho_1 \rho_2 - 1} \left(\frac{L_{t+1}^h}{L_t^h} \right)^{\rho_2 - \rho_1 \rho_2} (1 + r_{t+1}) \beta \exp \left(\theta \Delta Z_{t+1}^h + \Delta v_{t+1}^h \right) \right] = 1 \tag{11}$$

Results

The GMM point estimates of γ range from 1.18 to 1.39 across the four separable specifications tested; see Table 2. This equates to an estimated EIS of between 0.74 and 0.85. In the context of Euler equation-based studies, the GMM estimates are relatively precise with standard errors around between 0.17 and 0.24. The estimate for γ of 2.4 from the linearised Euler equation is quite a bit higher than the GMM estimates, but is very imprecise. This specific finding is very similar to that in Alan et al. (2009) using PSID data. Their simulation exercise in the same study highlights that estimates of γ based on linearised models are prone to severe biases in cases where the time dimension is less than 30 years or so (see also Attanasio and Low (2004)).

As noted above, β and σ_η^2 are not separately identifiable without further assumptions (or additional moment conditions). In models 4 and 5, we assume that measurement error accounts for 80 per cent of the overall variance in consumption growth, a proportion guided by our separate estimates in Table D.1. Given the variance in the Australian consumption data, this implies a measurement error variance of around 0.095 and yields estimates for β between 0.92 and 0.96 across two specifications.

The coefficient on change in household size gives an explanation for the ob-

Table 2: Euler equation estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Log. Lin.	EGMM	EGMM	EGMM	EGMM	EGMMn
Coeff. of relative risk aversion (γ)	2.381	1.364	1.285	1.389	1.18	2.245
	(1.261)	(0.166)	(0.202)	(0.236)	(0.199)	(0.734)
Discount factor (β)				0.921	0.962	
				(0.061)	(0.023)	
β/Meas'nt error in consumption		0.761	0.806			0.744
		(0.065)	(0.081)			(0.226)
Household size (θ)	0.524	0.358	0.329	0.376	0.294	1.445
	(0.28)	(0.069)	(0.077)	(0.093)	(0.068)	(1.115)
2008 dummy	-0.074		-0.07			
	(0.01)		(0.05)			
Constant	0.007					
	(0.004)					
Inflation measure for real interest rate	Survey	Survey	Survey	Survey	Bond M.	Survey
Instruments**:						
Lagged interest rate	Yes	Yes	Yes	Yes	Yes	Yes
Lagged consumption growth	No	Yes	Yes	No	Yes	Yes
Households	7225	6227	6227	7225	6227	3249
Observations	44966	35254	35254	44966	35254	16430

EGMM = exact GMM; EGMMn = exact GMM with non-separable preferences.
Standard errors in parentheses.
* Measurement error variance set to 80 per cent of variance in consumption growth, consistent with our findings in Appendix D.
** The set of instruments also includes a constant in all equations, as well change in HH size and the GFC dummy where present for columns (1) to (5). Lagged changes in household leisure time, and age and age-squared of the household head are also included for the model in column (6).

served hump-shaped profile of consumption over the life cycle, and is relevant for the construction of equivalence scales (see Gourinchas and Parker (2002) among others). Given our sample of households is restricted to couples who have stayed together, this coefficient basically measures the marginal utility associated with an extra child. The estimated coefficient averages about 0.3 for the GMM estimator, with a somewhat higher estimate of 0.48 using the linearised model. These estimates are in the range of those obtained in international Euler equation studies using food expenditure (for example, as Attanasio et al. (1999)), but somewhat below those in Alan et al. (2009).

In the non-separable model, we estimate a coefficient of relative risk aversion of 2.23 (column (6)). While somewhat higher than the separable model estimates, the much larger standard error indicates no statistically significant difference from the separable estimates at standard confidence levels.

We run a series of tests to check the sensitivity of our results to sample selection. We find that the exclusion of consumption growth outliers has very little impact on the results, nor does introducing a requirement that a respondent appears in at least five consecutive waves (as per Alan et al. (2009)). However, we find that including households whose head has begun or ended a marriage or de facto relationship during their time in the panel leads to implausibly low estimates of γ (well below 1) and instability in the GMM estimation procedure. In column 5, we experiment with an alternative real interest rate series, which is calculated using the inflation rate derived from market pricing of inflation-linked bonds. This leads to an only slightly lower estimate of γ.

Comparison with international estimates

Overall, our point estimates of γ for Australia based on the non-linear Euler equation are within the range of recent estimates obtained in the US.[23] In the most comparable US study, Alan et al. (2009) shows GMM estimates of γ ranging from 1.15 to 1.53 using data from the PSID, with a preferred point estimate of 1.45.[24] Our estimates of γ are also within the range of estimates found in Gourinchas and Parker (2002), who use the quite different approach of simulated method of moments. However, as noted earlier, these moderate estimates of γ contrast with the very high values needed to explain the large equity premium in the US (Kocherlakota, 1996).

6 Conclusion

Testing the CRRA assumption empirically is crucial as it underpins many theoretical and applied models of decision making under uncertainty, and our paper provides some qualified empirical support for the CRRA assumption for Australia. That said, given the difficulties associated with measurement error in the wealth data as well as the deeply endogenous relationship between risky asset share and wealth, further studies are needed to increase the level of confidence in this result. Such future studies will benefit from additional waves of panel data and may be able to approach the question from a different angle to ours, perhaps relying also on the subjective data in HILDA.

[23] We found very few recent studies outside the US.

[24] For comparison, their estimates of β ranged from 0.87 (which the authors considered implausibly low) to 0.99

While informative, our Euler estimates of the coefficient of relative risk aversion γ are also subject to uncertainty. Other approaches, such as the more structural simulated method of moments technique established by Gourinchas and Parker (2002), should be used to confront our results given the heavy reliance on this parameter in applied policy analysis.

A recent working paper by Iskhakov and Keane (2018) estimates a value of 0.8 for γ using simulated method of moments combined with HILDA data. This implies a much lower degree of relative risk aversion than we find.

Our estimates of γ also have interesting implications for precautionary saving that could be explored further.[25] In particular, Kimball (1990) shows that for consumers who face income (or other) uncertainty, the strength of the precautionary saving motive depends directly on their degree of prudence.[26] With CRRA preferences, our preferred estimates of γ imply a coefficient of relative prudence (given by $CU'''(C)/U''(C)$) between 2.2 and 2.4.[27] A structural life-cycle model could be used to explore what this implies for the proportion of household wealth attributable to precautionary saving. We are unaware of any Australian studies that have looked at this question.[28]

[25] There is a large empirical literature devoted to measuring the proportion of household wealth attributable to precautionary saving. Somewhat frustratingly, this literature has produced a very broad range of estimates. See Lusardi (1998) and Carroll and Samwick (1998) for examples at either extreme.

[26] Deaton (1991) and Carroll (2001b) among others point out that the prospect of a borrowing constraints can cause households to engage in precautionary saving even when the utility function does not exhibit prudence.

[27] As explained above, under CRRA preferences, the degree of relative prudence is simply equal to $\gamma + 1$.

[28] Hanel and Haisken-DeNew (2017), using HILDA data, find a large effect of uncertainty on savings (ie, a strong precautionary saving motive), but do not quantify its contribution to overall household wealth.

References

Ahmed, Naeem, Matthew Brzozowski, and Thomas Crossley, "Measurement Errors in Recall Food Consumption Data," *IFS Working Papers, Institute for Fiscal Studies (IFS)*, November 2006.

Alan, Sule, Kadir Atalay, and Thomas F. Crossley, "Euler Equation Estimation on Micro Data," *Macroeconomic Dynamics*, 2018, p. 1–26.

_ , Orazio Attanasio, and Martin Browning, "Estimating Euler equations with noisy data: two exact GMM estimators," *Journal of Applied Econometrics*, 2009, *24* (2), 309–324.

Altonji, Joseph G. and Lewis M. Segal, "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business and Economic Statistics*, 1996, *14* (3), 353–366.

Attanasio, Orazio P. and Hamish Low, "Estimating Euler Equations," *Review of Economic Dynamics*, April 2004, *7* (2), 405–435.

_ , James Banks, Costas Meghir, and Guglielmo Weber, "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics*, 1999, *17* (1), 22–35.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption Inequality and Partial Insurance," *American Economic Review*, December 2008, *98* (5), 1887–1921.

Bonin, Holger, Thomas Dohmen, Armin Falk, David Huffman, and Uwe Sunde, "Cross-sectional earnings risk and occupational sorting: The role of risk attitudes," *Labour Economics*, 2007, *14* (6), 926 – 937. Education and Risk.

Bound, John, Charles Brown, and Nancy Mathiowetz, "Chapter 59 - Measurement Error in Survey Data," in James J. Heckman and Edward Leamer, eds., *James J. Heckman and Edward Leamer, eds.*, Vol. 5 of *Handbook of Econometrics*, Elsevier, 2001, pp. 3705 – 3843.

Brzozowski, Matthew, Thomas F. Crossley, and Joachim K. Winter, "A comparison of recall and diary food expenditure data," *Food Policy*, 2017, *72*, 53 – 61. Special Issue: Food counts.

Cardak, Buly and Roger Wilkins, "The determinants of household risky asset holdings: Australian evidence on background risk and other factors," *Journal of Banking and Finance*, 2009, *33* (5), 850–860.

Carroll, Christopher, "Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation)," *The B.E. Journal of Macroeconomics*, 2001, *1* (1), 1–38.

Carroll, Christopher D., "A Theory of the Consumption Function, with and without Liquidity Constraints," *The Journal of Economic Perspectives*, 2001, *15* (3), 23–45.

_ **and Andrew A. Samwick**, "How Important is Precautionary Saving?," *The Review of Economics and Statistics*, 1998, *80* (3), 410–419.

Chiappori, Pierre-Andre and Monica Paiella, "Relative Risk Aversion Is Constant: Evidence From Panel Data," *Journal of the European Economic Association*, December 2011, *9* (6), 1021–1052.

Chioda, Laura, "Estimating Euler Equations with Measurement Error: A Nonparametric Approach," Available: https://pdfs.semanticscholar.org/f977/9ee90802d4b60fe76a466188602f6a737d73.pdf [15 November 2018], 2004.

Conlin, Andrew, Jouko Miettunen, Jukka Perttunen, Mikko Puhakka, and Rauli Svento, "Risk Aversion, Wealth and Personality," 2016. Unpublished working paper. Available at `https://www.univaasa.fi/en/research/groups/frg/researchseminars/financeseminars/file/`.

Deaton, Angus, "Saving and Liquidity Constraints," *Econometrica*, 1991, *59* (5), 1221–1248.

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption Over the Life Cycle," *Econometrica*, 2002, *70* (1).

Hall, Robert E., "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 1978, *86* (6), 971–987.

Hanel, Barbara and John Haisken-DeNew, "Keep Calm and Consume? Subjective Uncertainty and Precautionary Savings," July 2017. Melbourne Institute Working Paper No. 18/17.

Iskhakov, Fedor and Michael Keane, "Effects of Taxes and Safety Net Pensions on life-cycle Labor Supply, Savings and Human Capital: the Case of Australia," ANU Working Papers in Economics and Econometrics 2018-661, Australian National University, College of Business and Economics, School of Economics May 2018.

Kimball, Miles S., "Precautionary Saving in the Small and in the Large," *Econometrica*, 1990, *58* (1), 53–73.

Kocherlakota, Narayana, "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature*, 1996, *34* (1), 42–71.

Levy, Haim, "Absolute and relative risk aversion: An experimental study," *Journal of Risk and Uncertainty*, May 1994, *8* (3), 289–307.

Lusardi, Annamaria, "On the Importance of the Precautionary Saving Motive," *The American Economic Review*, 1998, *88* (2), 449–453.

Ogaki, Masao and Qiang Zhang, "Decreasing Relative Risk Aversion and Tests of Risk Sharing," *Econometrica*, 2001, *69* (2), 515–526.

Tsigos, Stamatios and Kevin Daly, ""Fair go" for all? Wealth and Risk Aversion of Australian Households*," *Australian Economic Papers*, 2016, *55* (3), 274–300.

Ventura, Eva, "A note on measurement error and Euler equations," *Economics Letters*, 1994, *45* (3), 305 – 308.

Wilkins, R. and C. Sun, "Assessing the quality of the expenditure data collected in the self-completion questionnaire," HILDA Project Discussion Paper Series no. 1/10, Melbourne Institute of Applied Economic and Social Research 2010.

Yagihashi, Takeshi and Juan Du, "Intertemporal elasticity of substitution and risk aversion: are they related empirically?," *Applied Economics*, March 2015, *47* (15), 1588–1605.

Appendix A Sample selection

Table A.1: Sample selection for the risky asset share regressions

	Dropped	Remaining
Unbalanced panel of respondents		317,738
Keep years for which there are wealth data	241,025	76,713
Keep working age	38,730	37,983
Keep household heads	17,280	20,703
Keep positive risky assets	1,250	19,453

Additional observations were lost in first differencing and because instrumental variables contained some missing values.

Table A.2: Sample selection for the Euler equation analysis

	Dropped	Remaining
Unbalanced panel of respondents		317,738
Keep couples who did not split up	36,423	281,315
Keep couples	151,319	129,996
Keep HH heads	54,347	75,649
Keep non credit constrained housheolds	36	75,613
Keep working age	1,717	73,896
Keep non-missing data	19,887	55,726

Additional observations were lost because of the presence of lagged variables in the equation and the instruments.

Appendix B Relative risk aversion estimates by net worth

Table B.1: Median relative risk aversion estimate based on risky asset share

Net worth quartile	Equities and shares	+Bond-like assets	+Bus. Equity and Super
Bottom	2.24	2.15	1.1
Second	2.6	2.44	1.1
Third	2.47	2.28	1.09
Top	2.43	2.01	1.08
Overall	2.47	2.18	1.09

Appendix C Impact of measurement error in wealth on risky asset regressions

Research on the impact of measurement error in household panel data has mainly focused on income and expenditure data (see Bound, Brown and Mathiowetz (2001) for a discussion), but it is likely that wealth data also suffers from substantial measurement error. This appendix explores the impact on the estimated relationship between risky asset share and wealth.

We begin by assuming that the logarithm of household wealth, $\log W_h^\star$, is normally distributed with mean μ_W and variance $\sigma^2_{W^\star}$, where the star denotes the true value of the variable.[29] Each household invests a share $\alpha^\star$ in risky assets $W_p^\star$ and $(1-\alpha^\star)$ in risk free assets $W_f^\star$.

Observed assets are subject to multiplicative measurement errors u_p and u_f, so that $W_j = W_j^\star u_j$ for $j = p, f$. We also assume that $\log(u_p)$ and $\log(u_f)$ are normally distributed with mean zero and variances σ_p^2 and σ_f^2 respectively. We further assume zero correlation between the two measurement errors.

For simplicity, we also assume that households have a common risky asset share $\alpha_h^\star = \bar{\alpha}^\star$. In this setting, cross-sectional variation in $W_p^\star$ and $W_f^\star$ are driven solely by a household's wealth level. Observed wealth, W, and the

[29] From now on we suppress the subscript h unless it is needed to remove ambiguity.

observed risky asset share, α, are then given by:

$$\begin{aligned} W &= W_p^\star u_p + W_f^\star u_f \\ &= \alpha^\star W^\star u_p + (1-\alpha^\star) W^\star u_f \\ \alpha &= \frac{\alpha^\star W^\star u_p}{\alpha^\star W^\star u_p + (1-\alpha^\star) W^\star u_f} \end{aligned} \tag{1}$$

We are interested in how measurement error of this kind affects OLS estimates of β_1 in the following regression:

$$\log(\alpha_h) = \beta_0 + \beta_1 \log(W_h) + \epsilon_h \tag{2}$$

In the absence of measurement error, OLS estimates of β_1 are unbiased and consistent and equal to zero given that the correlation between $\alpha_h^\star$ and $W_h^\star$ is zero by construction. However, in the presence of measurement error, the covariance between α_h and W_h becomes:

$$\begin{aligned} \mathrm{Cov}(\log \alpha_h, \log W_h) = \mathrm{Cov}[&\log(\alpha^\star W^\star u_p + (1-\alpha^\star) W^\star u_f), \log(\ \alpha^\star W^\star u_p) \\ &- \log(\alpha^\star W^\star u_p + (1-\alpha^\star) W^\star u_f)] \end{aligned} \tag{3}$$

This will generally not be equal to zero as measurement error affects both the dependent and independent variables. This fact also makes the direction of the bias ambiguous, as it will depend on the precise values of $\bar{\alpha}$ as well as the relative variances of the measurement errors. While there is no way of analytically quantifying this covariance expression, and therefore the direction and size of the bias, we can evaluate it numerically by simulating a large number of draws of $W^{\star}$, u_p and u_f.[30]

Figure C.1 shows how the bias in OLS estimates of β_1 varies with different values of $\bar{\alpha}$ based on the numerical evaluation of equation 3. The simulations assume that the variances of $\log(u_p)$ and $\log(u_f)$ are equal to each other and imposes a signal-to-noise ratio of four (ie, σ_r^2 and σ_f^2 are one-quarter the variance in $\log(W^{\star})$).[31] For values of α at or below 0.5, OLS estimates of β_1 have a negative bias that becomes more severe as $\bar{\alpha}$ gets smaller. For values of $\bar{\alpha}$ above 0.5, the bias in β_1 is positive but of relatively small magnitude. In short, the most severe (negative) bias in estimates of β_1 are likely to occur when the average risky share is less than 0.5, which is the situation we face in the main text.

[30] The covariance depends on non-linear transformations involving the sum of two log-normal random variables, which itself has no close form.

[31] The simulations involved 1,000,000 random draws of $W^{\star}$, u_p and u_f for each value of $\bar{\alpha}$.

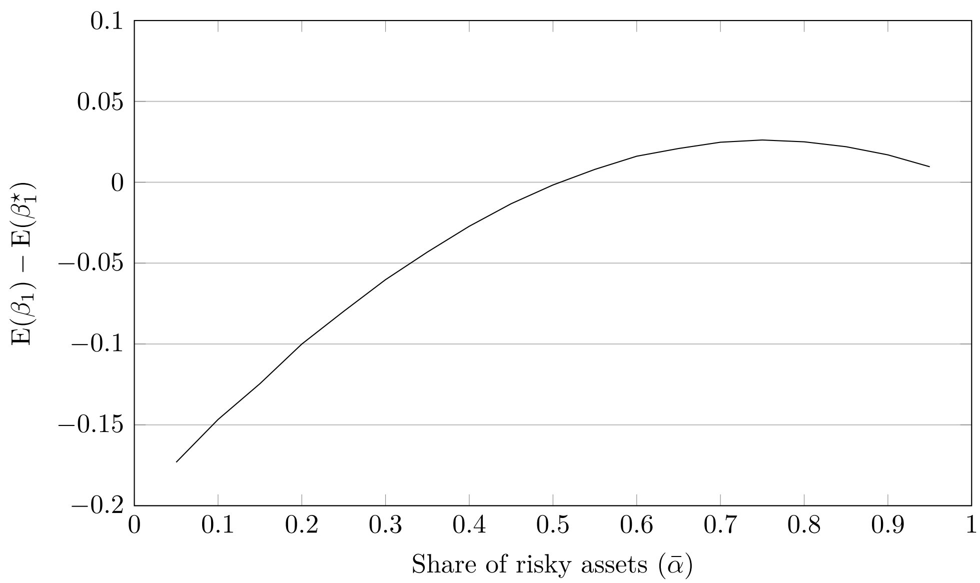

Figure C.1: Bias in OLS estimates of β_1

Appendix D Estimates of measurement error in grocery data

This appendix estimates the variance of measurement error in grocery expenditure data, which is an input into our Euler equation analysis in the main text.

We begin by making the common assumption that measurement error is multiplicative, so that $C_t^i = \eta_t^h C_t^\star$ where $C_t^\star$ is the true (unobserved) consumption for each household h.[32] We further assume that the measurement error η_t is an MA(1) process such that $\log \eta_t = \epsilon_t + \theta\epsilon_{t-1}$ where ϵ is a normally distributed iid error with with constant variance σ_ϵ^2. Consumption growth is then given by:

$$\Delta c_t = \Delta c_t^\star + \epsilon_t + (\theta - 1)\epsilon_{t-1} - \theta\epsilon_{t-2} \tag{1}$$

This implies the following variance and first- and second-order autocovariances for Δc_t:

$$\mathrm{Var}(\Delta c_t) = \mathrm{Var}(\Delta c_t^\star) + \sigma_\epsilon^2 + (1-\theta)^2\sigma_\epsilon^2 + \theta^2\sigma_\epsilon^2 \tag{2}$$

$$\mathrm{Cov}(\Delta c_t, \Delta c_{t-1}) = (\theta - 1)\sigma_\epsilon^2 - \theta(\theta - 1)\sigma_\epsilon^2 \tag{3}$$

$$\mathrm{Cov}(\Delta c_t, \Delta c_{t-2}) = -\theta\sigma_\epsilon^2 \tag{4}$$

where the autocovariance in true consumption growth one or more periods apart is assumed to be zero.[33] With three equations in three unknowns, the

[32] To avoid notational clutter, we suppress the h superscript from now on.

[33] This key identification assumption is equivalent to assuming that consumption is a

system is identified and can be estimated using generalized method moments (GMM) given at least four years of panel data.

We fit this model using data on self-reported grocery expenditure from HILDA's self completion questionnaire (SCQ) for the response of the household head and separately for the head's partner.[34] We also provide estimates model based on the spliced grocery and meals out series used in the main text, which covers a slightly longer period, from 2004 to 2016. In the latter case, the data is averaged across multiple responses from a single household where applicable.[35] Our GMM estimation is based on the identity weight matrix.[36]

Table D.1: GMM Estimates of measurement error in expenditure data

	(1)	(2)	(3)
Expenditure measure	Groceries	Groceries	Groc. & meals out
Data source	Head SCQ	Partner SCQ	Av. SCQ & HQ
Period	2006-2016	2006-2016	2004-2016
Var. of true cons. growth ($\text{Var}(\Delta c_t^\star)$)	0.0483	0.0593	0.0205
	(0.0158)	(0.0172)	(0.0018)
Variance of meas'nt error (σ_ϵ^2)	0.1001	0.0911	0.0699
	(0.0142)	(0.0158)	(0.0025)
MA(1) coefficient (θ)	0.0969	0.0356	0.0454
	(0.048)	(0.0889)	(0.015)
Proportion of $\text{Var}(\Delta c_t)$ due to Msmt. Error (%)	79	76	87
Num. of obs.	5502	5502	55554

Standard errors in parentheses.

martingale with drift, which is a common assumption in dynamic models of household consumption (see for example Blundell et al. (2008)).

[34] We restrict the sample to couple households where both members provide a *unique* response.

[35] The expenditure section of the SCQ can be completed by anyone in the household with "any responsibility for the payment of household bills...". Wilkins and Sun (2010) show that around one quarter of households provide more than one response each year.

[36] Altonji and Segal (1996) show that the identity weight matrix is superior to the 'optimal' weight matrix (the variance of the sample moment conditions) when estimating models of covariance structures in short panels. Also, with unbalanced panel data, there is no guarantee that the variance of the sample moment conditions will be positive semi-definite matrix, a problem that we encountered occasionally.

The estimates in Table D.1 imply that measurement error accounts for somewhere between 76 and 87 per cent of the variance in observed growth in grocery expenditure. This is within the range of international estimates, albeit towards the upper end. Interestingly, the process of averaging grocery responses across multiple household members, while reducing the noise, appears to reduce the signal even more so. This results in a higher estimated proportion of measurement error in this case (see column 3). The estimate of the MA(1) coefficient is significant in the first and last cases.