Essays on Consumption and Income Inequality In Ireland

by

Paul Joseph Morrin, BA (Mod. Econ), MSc

A thesis submitted in fulfilment of the requirements for the degree of Ph.D.

In the Department of Economics, National University of Ireland, Maynooth (NUIM), Maynooth, Co. Kildare, Ireland

May 2022

Head of Department: Dr. Tuvana Pastine

Department Supervisors: Professor Donal O'Neill (Lead Supervisor), Dr. Aedin Doris and Dr. Olive Sweetman

Maynooth University
National University of Ireland Maynooth

Table of Contents

List of Tables

List of Figures

Abstract

This thesis consists of essays examining the relationship between income and consumption in Ireland. If people with ‘rational expectations’ can freely borrow or save then ‘lifetime income’ or consumption is a more accurate measure of the distribution of household resources than current income. However, a number of issues such as measurement error in income and expenditure, formation of expectations and the relative importance of permanent and transitory shocks to income all pose challenges to anyone wishing to understand how households smooth consumption over income peaks and troughs to maintain living standards. I will examine some of these issues identified in the research in the Irish context through an analysis of household survey microdata.

Summary

This thesis consists of essays examining the relationship between income and consumption in Ireland. If people with 'rational expectations' can freely borrow or save then 'lifetime income' or consumption is a more accurate measure of the distribution of household resources than current income. However, a number of issues such as measurement error in income and expenditure, formation of expectations and the relative importance of permanent and transitory shocks to income all pose challenges to anyone wishing to understand how households smooth consumption over income peaks and troughs to maintain living standards. I will examine some of these issues identified in the research in the Irish context through an analysis of household survey microdata.

Chapter 3 covers measurement of income and expenditure. I find that income and expenditure are both well measured in the Irish data compared to the UK and US. I find that consumption has less inequality than the other measures and that the composition of those classified as poor changes depending on the outcome measure used.

Chapter 4 examines the 'retirement consumption puzzle' in the Irish context, which is a specific failure of the Permanent Income Hypothesis (PIH) in relation to retired consumers. I find that there is a fall in consumption at retirement age for Irish HBS respondents which is concentrated in services expenditures.

Chapter 5 looks at whether consumers underestimate their life expectancy. This analysis shows that household survey responses are well considered and reflect relevant socio-economic factors and health behaviours.

Chapter 6 examines earnings and income volatility in Ireland from 2005 to 2017 based on a unique linked dataset of administrative incomes to HBS household characteristics. I find that there was much higher volatility in Irish men's earnings during the "Great Recession" compared to levels reported in US and UK research since the 1970's. The main finding of my research is that households insured themselves against this volatility through male welfare and spousal incomes so that household income volatility was lower and flatter throughout the period when compared to male earnings.

Acknowledgements

Academic;

Donal, Olive and Aedin in NUI Maynooth for very patient and comprehensive support over the last five years, including during my 'gap year' in 2016/17 when I was figuring out if this was for me. I wasn't sure what to expect in 2015 when I started but the experience has been brilliant and I have been shown how professional economists think and express themselves. In addition, Dermot Leahy and David Madden for letting me attend their excellent courses.

Cormac O'Dea, Paul Redmond and Tim Callan who gave me ideas and more importantly encouragement at various times.

In CSO;

The Office for supporting me during my period of study.

My managers for their support – initially Jennifer Banim and then Pádraig Dalton.

Gerry Reilly and Claire Burke for additional HBS data for chapters 4 and 6; Berni Dunne for administrative data used in chapter 6; all provided in line with the CSO data protocol.

Gerard Keogh, Patrick Quill, Tim Linehan, Brian King and Brian Stanley for very helpful and good humoured technical discussions from time to time.

John Dunne, Eoin Flaherty and Steve MacFeely for moral support and encouragement.

At home;

My wife for giving me time in the evenings and weekends to get through the workload and for support and encouragement generally.

Issey and Chloe Tighe for family support in the last year when it was much needed.

My children Niall, Sé and Ann who arrived during chapters 3, 4 and final drafting and offered support and inspiration in their own way.

1. Introduction

The relationship between income and consumption has been central to the development of economic thought over the past century. This is not just a theoretical argument, since policymakers' beliefs about this relationship have informed decisions about austerity and government intervention in the "Great Recession". These decisions had implications for the living standards of citizens. Classical economics argues that people are rational and not subject to huge swings in confidence. If people with 'rational expectations' can freely borrow or save then 'lifetime income' or consumption is a more accurate measure of the distribution of household resources than current income. Friedman (1957) introduced the 'Permanent Income Hypothesis' (PIH) to explain how households smooth consumption over income peaks and troughs to maintain living standards, while Modigliani (1966) refined this analysis to reflect rising and falling consumption needs based on age and family composition with the 'Life Cycle Hypothesis'. However, changes in income and consumption over time and across age cohorts tend to be closely related in statistical data from household surveys or the National Accounts.

Various failures of the PIH have been examined in economic research, including restrictions on borrowing and failures of rational expectations. Browning and Crossley (2001) outlined how more sophisticated models which look at consumption choices jointly with other choices, durability of goods and portfolios of assets reveal more information about income smoothing. Heterogeneity also needs to be allowed for in these models. Measurement problems with income and consumption have also been examined to ensure that analysis of their relationship is well based. There are also issues with using single cross section income as a predictor of lifetime income and income risk. Aaberge and Mogstad (2015) used full career histories of male Norwegian workers to examine 'life-cycle bias', which is the extent to which current incomes are representative of lifetime incomes. This is likely to be an issue in Irish cross-section analysis given the increase in living standards experienced between the two crises. There is also extensive analysis of whether outcomes related to volatility or variance of incomes are consistent with the PIH, which assumes consumers can distinguish between transitory and permanent income shifts. Blundell et al. (2008) found some partial insurance for permanent income shocks in consumption and almost complete insurance of transitory shocks, though low-wealth households were more exposed to transitory risk. I will examine some of the issues with the PIH identified in the research in the Irish context through analysis of household survey microdata.

Chapter 2 summarises the main theoretical models used to analyse consumption over the life-cycle. I begin with a life-cycle model with no uncertainty, initially in general form and subsequently with Constant Relative Risk Aversion and quadratic utility functions specified. These specifications allow 'closed form' solutions to be identified for utility problems. I then add uncertainty about incomes to the quadratic utility model which results in Friedman's PIH in expected incomes. I then look at how failures of the PIH can be accommodated in utility problems. Firstly, I look at precautionary savings with convex marginal utility, where consumers prefer less income risk to more for a given expected level of income. I then add borrowing constraints and the 'buffer stock' model (where consumers aim to achieve a minimum level of liquid assets) to precautionary savings which better reflect the financial conditions facing consumers.

Chapter 3 covers measurement of income and expenditure in the Household Budget Survey (HBS), which includes both measures for the same respondent households. The HBS is a representative sample survey of Irish households that has been conducted every five years since 1994/1995, and less frequently before then. It includes detailed data on expenditure, income and household characteristics. The main purpose of the survey at the moment is to supply updated weights for the Consumer Price Index. The HBS will be moving to an annual cycle in the coming years, which will greatly enhance its potential for monitoring poverty and inequality. I find that income and expenditure are both well measured in the survey, and particularly that low incomes are accurate compared to the UK and US equivalent surveys. The reason for good measurement of incomes may be the relatively simple Irish benefits system. I examine this using waves in the years before and after a change in the survey process. I also examine the distributions of income, expenditure and consumption. I find that consumption has less inequality than the other measures, which is in line with the PIH. I also find that a group who experienced a transitory income shock in the "Great Recession" replaced a group with permanently low incomes at the bottom of the income distribution, while maintaining expenditure levels in line with the PIH. My main contributions to the literature in this chapter is to illustrate the much lower measurement error for Irish low incomes than in similar UK and US surveys and highlight the different profiles of poverty and inequality for consumption compared to income in Ireland.

Chapter 4 examines the 'retirement consumption puzzle' in the Irish context, which is a specific failure of the PIH in relation to retired consumers. International research shows that these consumers significantly reduce their consumption in retirement, which is sometimes presented as a failure of rational expectations, since no news about incomes is revealed to these consumers in retirement. There are different findings on this point based on different approaches to measurement, which inform my analysis. I find that the drop in consumption at retirement age for Irish HBS respondents is concentrated in services expenditures. This may be a failure of the PIH since these expenses are not particularly related to work or child rearing. My main contributions to the literature in this chapter are to show the concentration of reduced consumption in services and highlight the importance of housing services flows when considering Irish consumption.

Chapter 5 looks at whether consumers underestimate their life expectancy, which is one potential reason for the low level of savings (income less expenditure) in the HBS for working age households. Retirement is expected to be funded by savings during working age in the Life-Cycle Hypothesis. I analyse self-assessed survival expectations from the TILDA survey using approaches from international research. This analysis shows that the responses are well considered and reflect relevant socio-economic factors and health behaviours. My main contribution to the research is that a group who state probability of survival to future ages at 100%, who are treated as outliers or incorrect responses in previous research, have good predictive value in the Irish data. Apparent inconsistencies with Official estimates of life expectancy on an age and sex basis are substantively reduced when these responses are included in the comparison. My main contributions to the literature in this chapter are to show the importance of using unrestricted samples in this type of analysis and the relative lack of bias in the Irish responses compared to other countries, which may have policy implications.

Chapter 6 examines earnings and income volatility in Ireland from 2005 to 2017 based on a unique linked dataset of administrative incomes to HBS household characteristics. I find that there was much higher volatility in Irish men's earnings during the "Great Recession" compared to levels reported in US and UK research since the 1970's. This increased volatility was concentrated in low wage and low education employments. Men's earnings recovered quickly in 2013/14 as employment, earnings, and hours worked improved simultaneously. The main finding of my research is that households insured themselves against this volatility through male welfare and spousal incomes, so that household income volatility was lower and

flatter throughout the period. My main contributions to the literature in this chapter is to show the different pattern of income volatility and partial insurance in Ireland compared to comparable international studies and highlight the additional information on income dynamics available from panel data compared to cross section survey data. This is the first time such an analysis has been conducted for Ireland.

There are a number of common themes in the four empirical chapters. Measurement issues are addressed throughout, as good quality data are needed to draw accurate conclusions in subsequent analyses. Formation of expectations is examined in two of the four chapters. Finally, application of new methodologies in the Irish context, such as consumption and volatility, reveals different patterns to cross section analysis of income. Some directions for future research are outlined in Chapter 7.

2. Literature Review

In this chapter I review the main theoretical models used to analyse consumption over the life-cycle. I begin with a life-cycle model with no uncertainty, initially in general form and subsequently with Constant Relative Risk Aversion and quadratic utility functions specified. These specifications allow 'closed form' solutions to be identified. I then add uncertainty about incomes to the quadratic utility model which results in Friedman's Permanent Income Hypothesis (PIH) in expected incomes. I then look at how failures of the PIH can be accommodated. Firstly, I look at precautionary savings with convex marginal utility, where consumers prefer less income risk to more for a given expected level of income. I then add borrowing constraints and consider the 'buffer stock' model (where consumers aim to achieve a minimum level of liquid assets) to precautionary savings which better reflect the financial conditions facing consumers.

2.1 Life-cycle model with no uncertainty

2.1.1 In general

A general form for the life-cycle model of consumption under certainty is set out in this section. A single person household lives T periods from t=0 to t=T-1 with lifetime preferences defined over consumption sequences $\{c_t\}_{t=0}^{T-1}$:

$$U = \sum_{t=0}^{T-1} \beta^t u(c_t) \qquad (1)$$

Where $0<\beta<1$ is the discount factor, c_t is the household's consumption in period t and $u(c_t)$ measures the utility the household derives from consuming c_t in period t. $u(c_t)$ is assumed to satisfy the following conditions:

$u'(c_t) > 0$ and

$u''(c_t) < 0$.

Utility increases with additional consumption, as shown by the positive first derivative, while it increases at a decreasing rate based on the negative second derivative. These conditions reflect a concave utility function, where the household is risk averse (i.e., would give up a higher expected present value of future utility for certainty today). Such a utility function is illustrated in Figure 2.1:

Figure 2.1 Concave utility function

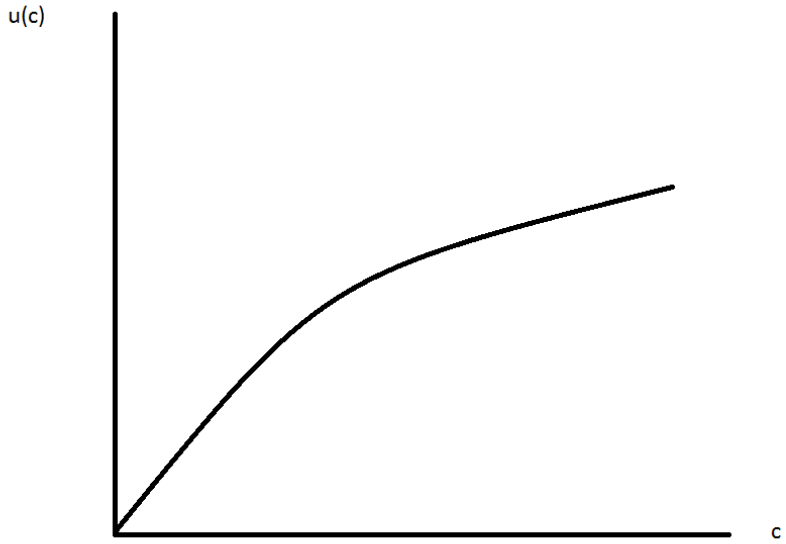

There are a number of assumptions underlying this representation of preferences;

- Certainty – in particular, there is certainty about *T* and household composition.
- Aggregation - aggregate consumption represents expenditures over many commodities - $c_t = \sum_i p_{i,t} * q_{i,t}$ over commodities *i* with prices and quantities *p* and *q*. It may not be possible to infer consumer behaviour or preferences from aggregate consumption data since the consumer's underlying preferences are hidden by aggregation. If preferences are homothetic[1] over commodities, then preferences in form ***u(q,p)*** can be 'aggregated' into an expression $u(c)$ where c=**p.q.**
- Separability – the life-cycle model presented here is based on optimal consumption, while other arguments enter utility in practice like labour supply or leisure. For this case, I assume that preferences are separable over different arguments and can be written as $f(c,z) = u(c) + v(z)$. There are no 'cross' effects of one argument over the marginal effect of the other on the function f, i.e. the cross partials are zero. Heckman (1974) proposed that consumption and labour supply are not separable since there are costs related to working which increase consumption during employment.
- Time additivity –it is assumed that the marginal utility of consumption at *t* only depends on consumption at that time. There is a wide range of literature on deviations from this

[1] This means that rich and poor consumers are equally averse to proportional changes in consumption.

assumption, including durability analysis for goods, habit formation and habit formation based on the consumption of other households.

- Intertemporal Marginal Rate of Substitution (IMRS) – the IMRS looking forward from t to t+1 is $\beta^{t+1}u'(c_{t+1})/\beta^{t}u'(c_t)$ and from time t+1 to t+2 is $\beta^{t+2}u'(c_{t+2})/\beta^{t+1}u'(c_{t+1})$ – i.e., both are equal given exponential discounting. The IMRS between consumption at any two points in time only depends on the distance between the points. Exponential discounting does not allow preferences to become time-inconsistent. As a result, households won't want to re-optimise over time.

The budget constraint of the household is based on the accumulation of assets (a_{t+1}), where a_{t+1} is the amount of assets at the beginning of t+1. Savings, given as income less consumption ($y_t - c_t$) is invested at a certain rate of interest r, which is assumed to be paid at the start of period t+1. Therefore:

$$a_{t+1} = (1+r)(a_t + y_t - c_t)$$

This assumption understates the amount of interest paid over the life cycle compared to a continuous compounding assumption since interest accumulates within periods. I derive the intertemporal budget constraint for the household initially:

$$a_0 = (1+r)^{-1}(a_1) - y_0 + c_0$$

By substituting and solving forward repeatedly for a_t I get:

$$a_0 = \sum_{t=0}^{T-1}(1+r)^{-t}(c_t - y_t) + (1+r)^{-T}a_T$$

The household cannot die in debt , i.e. $a_T \geq 0$, which is the basis for the constraint on the lifetime budget:

$$\sum_{t=0}^{T-1}(1+r)^{-t}(c_t) \leq a_0 + \sum_{t=0}^{T-1}(1+r)^{-t}(y_t)$$

Thus the present value of consumption is initial wealth plus the present value of income $\sum_{t=0}^{T-1}(1+r)^{-t}(y_t)$.

The solution for optimal consumption under certainty is obtained by maximising the objective function subject to the lifetime budget constraint **(2):**

$$\max_{\{c_t\}_{t=0}^{T-1}} \sum_{t=0}^{T-1} \beta^t u(c_t)$$

s.t.

$$\sum_{t=0}^{T-1} (1+r)^{-t}(c_t) \leq a_0 + \sum_{t=0}^{T-1} (1+r)^{-t}(y_t) \qquad (2)$$

I solve this problem by setting up the Lagrangian, where λ>0 is the Lagrange multiplier on the intertemporal budget constraint, assuming that the constraint is binding:

$$L = \sum_{t-0}^{T-1} \beta^t u(c_t) + \lambda \left(a_0 + \sum_{t=0}^{T-1} (1+r)^{-t}(y_t) - \sum_{t=0}^{T-1} (1+r)^{-t}(c_t) \right) \quad (3)$$

The first order condition for c_t is:

$$\frac{\partial L}{\partial c_t} = \beta^t u'(c_t) - \lambda(1+r)^{-t} = 0$$

So:

$$\lambda = (1+r)^t \beta^t u'(c_t)$$

The first order condition holds for each period under the assumptions described earlier, so for period t+1:

$$\frac{\partial L}{\partial c_{t+1}} = \beta^{t+1} u'(c_{t+1}) - \lambda(1+r)^{-(t+1)} = 0$$

So:

$$\lambda = (1+r)^{t+1} \beta^{t+1} u'(c_{t+1})$$

These conditions set up the intertemporal Euler equation, which is the 'law of motion' for optimal consumption in the life-cycle model;

$$(1+r)^t \beta^t u'(c_t) = (1+r)^{t+1} \beta^{t+1} u'(c_{t+1})$$

Solving the above gives:

$$\frac{u'(c_t)}{u'(c_{t+1})} = \frac{(1+r)^{t+1} \beta^{t+1}}{(1+r)^t \beta^t}$$

$$\frac{u'(c_t)}{u'(c_{t+1})} = \beta(1+r)$$

$\beta=\frac{1}{(1+\rho)}$, where ρ is the discount or time preference rate. Where ρ is positive, the consumer is impatient, and prefers consumption now to consumption tomorrow. The intertemporal Euler equation can be restated as:

$$u'(c_t) = \left(\frac{1+r}{1+\rho}\right) u'(c_{t+1})$$

If $r>\rho$ then $u'(c_t) > u'(c_{t+1})$, so $c_t < c_{t+1}$ since utility is concave. Consumption will grow over time since the consumer has an incentive to save towards future consumption.

2.1.2 Life-cycle model with Constant Relative Risk Aversion (CRRA) utility

A utility function needs to be specified in order to get a 'closed form' solution for the general form problem. The CRRA utility function is one possible specification. The objective function based on CRRA utility is:

$$U = \sum_{t-0}^{T-1} \beta^t \left(\frac{{c_t}^{1-\theta}}{1-\theta}\right) \qquad (4)$$

Where θ is the CRRA coefficient. The first and second derivatives of the CRRA function are;

$$U'(c_t) = \beta^t c_t^{-\theta} > 0$$

$$U''(c_t) = -\beta^t \theta c_t^{-\theta-1} < 0$$

i.e. Marginal Utility is always positive and diminishing. The shape of the utility function is given in Figure 2.2 for alternative values of the CRRA coefficient θ:

Figure 2.2 Constant Relative Risk Aversion utility function

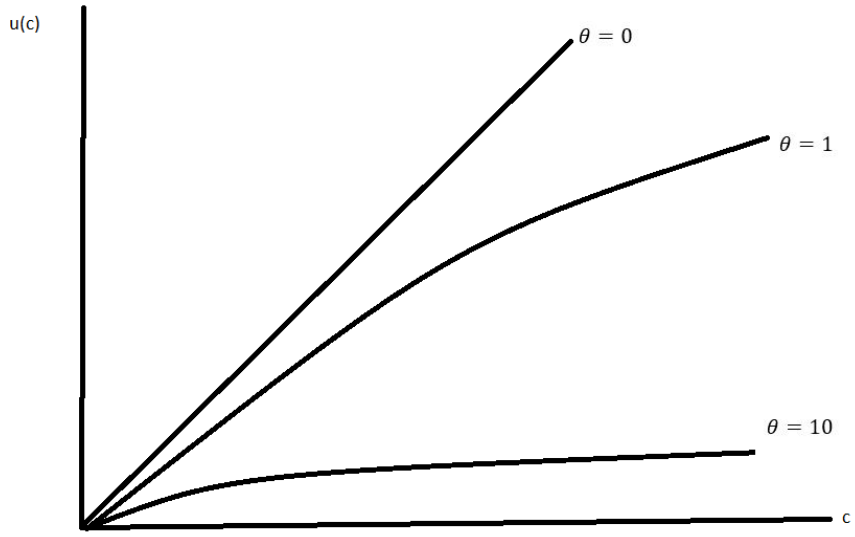

The CRRA utility function has the desirable property that the consumer's decision-making process is not affected by scale, so the optimal fraction of income consumed is unaffected by the consumer's level of income. However, there are some limitations to using the CRRA function in practice, since the risk aversion parameter also determines the intertemporal elasticity of substitution ($1/\theta$) and the consumer's prudence level. Risk aversion captures an individual's dislike of risk, the intertemporal elasticity of substitution is the consumer's preference for consumption over time, while prudence refers to how the disutility associated with additional risk varies with wealth. These are all distinct features of an individual's preference yet with the CRRA are all determined by the same parameter.

As with the general form, I use the lifetime budget constraint under certainty **(2)** with the CRRA objective function **(4)** to set up the Lagrangian:

$$L = \sum_{t-0}^{T-1} \beta^t \left(\frac{{c_t}^{1-\theta}}{1-\theta} \right) + \lambda \left(a_0 + \sum_{t=0}^{T-1} (1+r)^{-t} (y_t) - \sum_{t=0}^{T-1} (1+r)^{-t} (c_t) \right)$$

The first order condition for c_t is:

$$\frac{\partial L}{\partial c_t} = \beta^t (c_t)^{-\theta} - \lambda (1+r)^{-t} = 0$$

Which implies:

$$\lambda = (1+r)^t \beta^t (c_t)^{-\theta}$$

The first order condition holds for each period under the assumptions described earlier, so for period t+1:

$$\frac{\partial L}{\partial c_{t+1}} = \beta^{t+1} (c_{t+1})^{-\theta} - \lambda(1+r)^{-(t+1)} = 0$$

And:

$$\lambda = (1+r)^{t+1} \beta^{t+1} (c_{t+1})^{-\theta}$$

These conditions set up the intertemporal Euler equation:

$$(1+r)^t \beta^t (c_t)^{-\theta} = (1+r)^{t+1} \beta^{t+1} (c_{t+1})^{-\theta}$$

And:

$$\frac{(c_t)^{-\theta}}{(c_{t+1})^{-\theta}} = \beta(1+r)$$

$$\frac{(c_t)}{(c_{t+1})} = [\beta(1+r)]^{-1/\theta}$$

This differs from the intertemporal Euler equation in the general case since the time preference for consumption is determined by the CRRA coefficient θ as well as the discount and interest rates. When $\theta > 0$, the utility function is concave. Consumers will reduce current consumption to keep the marginal utility of consumption constant over time when interest rates increase. A more positive level of $\theta > 1$ indicates a stronger response to interest rates, which will further reduce current consumption.

2.1.3 Life-cycle model with quadratic utility

The quadratic utility function is another possible specification which obtains a closed form solution to the consumer's problem based on the objective function:

$$U = \sum_{t-0}^{T-1} \beta^t \left(\alpha c_t - \gamma \frac{{c_t}^2}{2} \right) \qquad (5)$$

Where $\gamma, \alpha > 0$.

The first and second derivatives of the quadratic function are:

$$U'(c_t) = \beta^t(\alpha - \gamma c_t)$$

$$U''(c_t) = -\beta^t\gamma$$

Utility can increase or decrease with consumption depending on the difference in the coefficients and marginal utility is a linear function of c_t. Where $\alpha - \gamma c_t < 0$, utility is decreasing with c_t. In this case, risk aversion represented by γ is greater than the fixed point α. The shape of the utility function is given in Figure 2.3 below;

Figure 2.3 Quadratic utility function

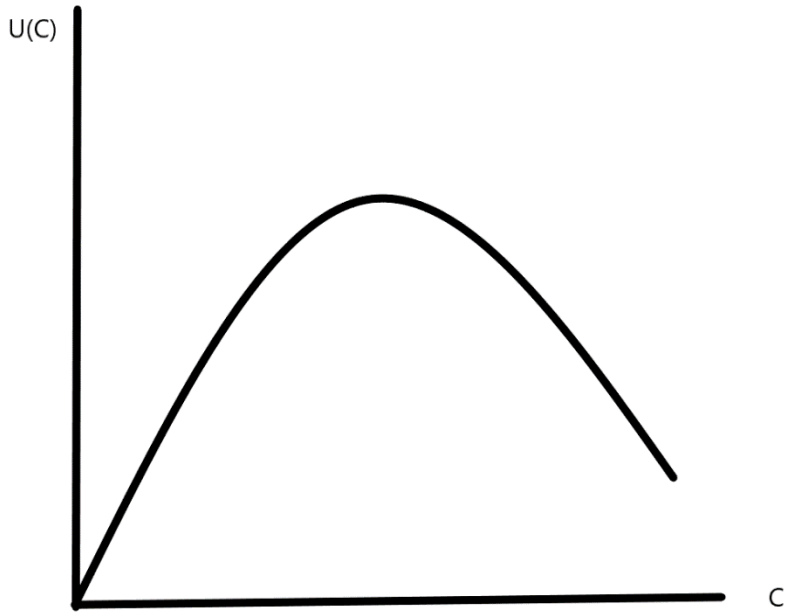

Quadratic preferences have undesirable features, as they imply a 'bliss point' for consumption and possible negative consumption which are implausible in reality. However, quadratic preferences are useful in certain circumstances, and this is discussed below.

As with the general form, I use the lifetime budget constraint under certainty **(2)** with the quadratic utility objective function **(5)** to set up the Lagrangian;

$$L = \sum_{t-0}^{T-1} \beta^t \left(\alpha c_t - \gamma \frac{c_t^2}{2}\right) + \lambda \left(a_0 + \sum_{t=0}^{T-1}(1+r)^{-t}(y_t) - \sum_{t=0}^{T-1}(1+r)^{-t}(c_t)\right)$$

The first order condition for c_t is:

$$\frac{\partial L}{\partial c_t} = \beta^t(\alpha - \gamma c_t) \quad - \lambda(1+r)^{-t} = 0$$

So:

$$\lambda = (1+r)^t \beta^t (\alpha - \gamma c_t)$$

The first order condition holds for each period under the assumptions described earlier, so for period t+1:

$$\frac{\partial L}{\partial c_{t+1}} = \beta^{t+1} (\alpha - \gamma c_{t+1}) \quad - \lambda (1+r)^{-(t+1)} = 0$$

And:

$$\lambda = (1+r)^{t+1} \beta^{t+1} (\alpha - \gamma c_{t+1})$$

These conditions determine the intertemporal Euler equation:

$$(1+r)^t \beta^t (\alpha - \gamma c_t) = (1+r)^{t+1} \beta^{t+1} (\alpha - \gamma c_{t+1})$$

$$\frac{(\alpha - \gamma c_t)}{(\alpha - \gamma c_{t+1})} = \beta(1+r)$$

With the additional assumption that $\beta(1+r) = 1$, consumption becomes constant over time or $c_t = \bar{c}$. When consumers can freely save and borrow to maintain consumption this assumption is more realistic. Substituting this into the intertemporal budget constraint with $a_T = 0$ instead of $a_T \geq 0$ and using the geometric sum formula for a finite series and the assumption that $\beta(1+r) = 1$:

$$\sum_{t=0}^{T-1} (1+r)^{-t} (\bar{c}) = \frac{1-\beta^T}{1-\beta} \bar{c} = a_0 + \sum_{t=0}^{T-1} (1+r)^{-t} (y_t)$$

This gives the life-cycle version of Friedman's Permanent Income Hypothesis (1957) – individual consumption at a point in time *t* is determined by lifetime resources and not by income in period *t.* Rearranging:

$$\bar{c} = \frac{1-\beta}{1-\beta^T} \left(a_0 + \sum_{t=0}^{T-1} (1+r)^{-t} (y_t) \right)$$

Lifetime resources are determined by initial assets a_0 and lifetime income y_t. The fraction of lifetime resources consumed depends on the discount rate β and the length of time *T* over which the consumer is optimising. The $\frac{1-\beta}{1-\beta^T}$ term converts discounted lifetime resources into an

expected annuity value which is 'permanent income'. This implies that consumption should not respond much to transitory changes in income, since these will not have a strong effect on permanent income.

In summary, the main characteristics of certainty-based models are;

- Consumption does not track current income – as lifetime resources are known with certainty, the marginal propensity to consume depends on the relationship between the discount rate and the interest rate, which is simplified to an annuity rate in Friedman's model. If $r>\rho$ consumption will grow over time due to savings, while if $r<\rho$ the consumer is impatient and consumption will fall over time as in earlier models;
- Permanent income is an intermediate stage in the calculation of consumption where consumers calculate the annuity value of lifetime resources. Transitory shocks add less to lifetime resources, but the marginal propensity to consume is not affected by whether shocks are permanent or transitory. All shocks are predicted so unanticipated shocks are not a factor in consumption planning;
- Consumption at retirement is planned in advance through accumulation of wealth and subsequent dissaving. The consumer is also certain about remaining life years which facilitates maintenance of consumption rates in retirement.

2.2 Life-cycle with uncertainty and quadratic utility

Using the quadratic model with uncertainty, expectations enter the utility objective function:

$$E(U) = E_0\left[\sum_{t-0}^{T-1} \beta^t \left(\alpha c_t - \gamma \frac{{c_t}^2}{2}\right)\right] \qquad (6)$$

As with the general form, I use the lifetime budget constraint under certainty **(2)** with the quadratic utility objective function with uncertainty **(6)** to set up the Lagrangian:

$$L = E_0\left[\sum_{t-0}^{T-1} \beta^t \left(\alpha c_t - \gamma \frac{c_t^2}{2}\right)\right] + \lambda\left(a_0 + \sum_{t=0}^{T-1}(1+r)^{-t}(y_t) - \sum_{t=0}^{T-1}(1+r)^{-t}(c_t)\right)$$

The first order condition for choice of consumption in any future period in time 1 is:

$$\frac{\partial L}{\partial c_t} = E_0\left[\beta^t(\alpha - \gamma c_t)\right] - \lambda(1+r)^{-t} = 0$$

Taking β^toutside the expectation:

$$\lambda = (1+r)^t \quad \beta^t E_0[(\alpha - \gamma c_t)]$$

And with the additional assumption that $\beta(1+r) = 1$:

$$\lambda = E_0[(\alpha - \gamma c_t)]$$

Taking the scalars α and γ outside the expectation:

$$\lambda = \alpha - \gamma E_0(c_t)$$

Also, since first period consumption is known at time 0, then the first order condition for period 0 is:

$$\lambda = \alpha - \gamma c_0$$

Combining these last two expressions for λ I get:

$$E_0(c_t) = c_0$$

This condition holds for each consumption value chosen by the household at time 1 and also in general, so that:

$$E_{t-1}(c_t) = c_{t-1}$$

Also;

$$E_{t-1}(c_{t+1}) = c_{t-1} \; \forall \; t \geq 1$$

If the lifetime budget constraint is binding, then it is also binding in expectations:

$$\sum_{t}^{\infty}(1+r)^{-t}E_t(c_t) = \frac{1-\beta^T}{1-\beta}c_{t-1} = a_{t-1} + \sum_{t}^{\infty}(1+r)^{-t}E_t(y_t)$$

Which gives Friedman's Permanent Income Hypothesis (1957) in expected incomes:

$$c_{t-1} = \frac{1-\beta}{1-\beta^T}\left(a_{t-1} + \sum_{t}^{\infty}(1+r)^{-t}E_t(y_t) \quad \right)$$

This presentation assumes a long but finite planning horizon *T*. If consumers apply the same discount rate to bequests as for consumption, then outcomes are similar based on intergenerational models. Also, the discount rate is now a variable instead of an annuity value as it can change over time. Consumers reoptimise their consumption planning over time in

response to 'news' that arrives with uncertainty. In summary, the main characteristics of life cycle models with uncertainty and quadratic preferences are:

- Consumers behave as if future income is certain in these models, hence they are known as 'certainty equivalent'. Consumers react to 'news' about their incomes by freely borrowing and lending to smooth their consumption over income shocks, and consumption does not track short term incomes;
- Another feature of these models is that consumers have a much lower marginal propensity to consume out of transitory income than permanent income shocks, and it is assumed that consumers can distinguish these shocks;
- Predicted shocks are already included in expectations and permanent income, and consumers borrow to maintain their consumption levels. Only unpredicted shocks affect consumption;
- 'News' does not affect income in retirement, as it only relates to income earned during working age so no news is revealed on retirement. Consumption in retirement is funded by running down assets.

2.3 Precautionary savings with convex marginal utility[2]

When an increase in future risk increases the level of current savings, the additional saving is referred to as precautionary savings. To simplify the presentation of precautionary savings, the general utility problem can be expressed as an 'income fluctuations problem' based on the level of assets held between consecutive periods. The level of assets held can then be used as a constraint in consumption planning, which is more realistic for low income or low assets consumers in particular than the borrowing and lending assumptions used so far:

$$E(U) = E_t \left[\max_{\{c_t\}_{t=0}^{T-1}} \sum_{t=0}^{T-1} \beta^t u(c_t) \right]$$

s.t.

$$a_{t+1} = R(a_t + y_t - c_t)$$

Where (R = 1+*r*).

[2] This section draws on lecture notes provided online by Isaac Baley https://www.isaacbaley.com/advanced-macro-ii.html.

The constraint can be further simplified in terms of 'cash in hand'[3] x_t, which is the amount of liquid resources the household has access to at the beginning of each period:

$$x_t = a_t + y_t$$

So that the constraint of the problem becomes**:**

$$x_{t+1} = R(x_t - c_t) + y_{t+1} \qquad (7)$$

The problem can be written in Bellman equation form as an objective function:

$$V_t(x_t) = \max_{\{c_t\}} u(c_t) + \beta E_t [V_{t+1}(x_{t+1})] \qquad (8)$$

s.t.

$$x_{t+1} = R(x_t - c_t) + y_{t+1}$$

Substituting the budget constraint into the objective function gives a maximisation problem:

$$V_t(x_t) = \max_{\{c_t\}} u(c_t) + \beta E_t[V_{t+1}(R(x_t - c_t) + y_{t+1})] \quad (9)$$

The f.o.c. is:

$$\frac{\partial V_t(x_t)}{\partial c_t} = u'(c_t) + \beta E_t \left[V'_{t+1}(x_{t+1}) \left(\frac{\partial (R(x_t - c_t) + y_{t+1})}{\partial c_t}\right)\right] = 0$$

$$\frac{\partial V_t(x_t)}{\partial c_t} = u'(c_t) + \beta E_t[V'_{t+1}(x_{t+1})(-R)] = 0$$

Taking the scalar R outside the expectations operator to give the f.o.c.:

$$u'(c_t) = R \beta E_t[V'_{t+1}(x_{t+1})]$$

The envelope theorem implies that the total derivative of a value function with respect to any choice variable will be zero for optimising consumers, as the f.o.c. always holds. This allows c_t to be treated as a constant, i.e. $\frac{\partial c_t^*}{\partial x_t} = 0$. Taking the derivative of the Bellman equation (9) with respect to x_t gives:

$$V'_t(x_t) = R \beta E_t[V'_{t+1}(x_{t+1})]$$

[3] With the additional assumption that income shocks are independent and identically distributed so that assets and incomes always enter additively and incomes are not auto-correlated

Substitute this expression into the f.o.c. to get:

$$u'(c_t) = V'_t(x_t)$$

And rolling on one period:

$$u'(c_{t+1}) = V'_t(x_{t+1})$$

Substituting in the f.o.c. gives the Euler equation for consumption:

$$u'(c_t) = R\,\beta\,E_t[u'(c_{t+1})]$$

To examine this model in more detail, we need to distinguish between risk-aversion and prudence when describing an individual's preference. A risk averse individual is averse to mean preserving spreads in outcomes or equivalently the utility premium, which measures the change in utility from additional risk, is negative, such that a consumer suffers pain when risk is added to wealth. This is equivalent to specifying that the underlying utility function is concave ($U'' < 0$).

Prudence on the other hand considers how the utility premium changes as wealth changes. An individual is said to be prudent whenever the utility premium decreases with wealth. Under the assumption of risk aversion, whereby the utility premium is negative, the pain associated with additional risk is smaller the greater is wealth. It can be shown that prudence is equivalent to $U''' > 0$, or that marginal utility is convex.

The shape of the convex marginal utility function is given in Figure 2.4 below:

Figure 2.4 Convex marginal utility function

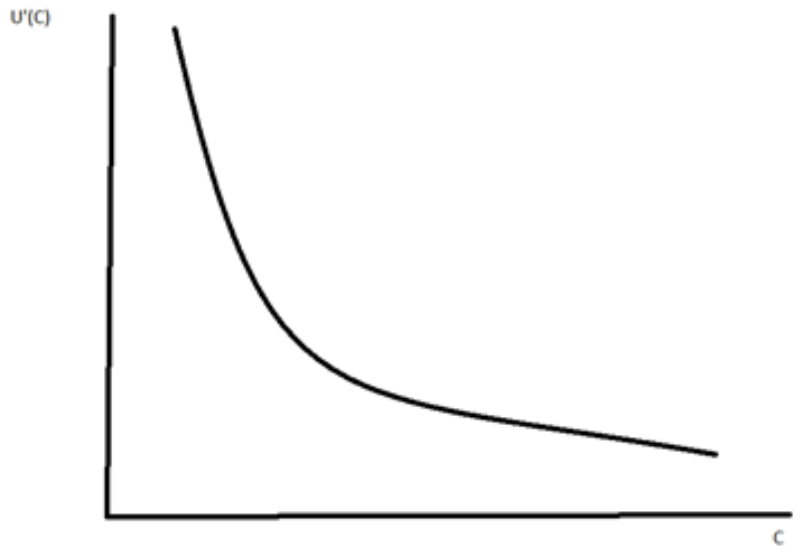

Jensen's inequality for convex functions, f, implies:

$$f(E[X]) \leq E[f(X)]$$

The budget constraint **(7)** can be re-stated in terms of c_t:

$$x_{t+1} = Rx_t - Rc_t + y_{t+1}$$

$$c_t = x_t + \frac{y_{t+1} - x_{t+1}}{R}$$

Substituting back in to the f.o.c.:

$$u'\left(x_t + \frac{y_{t+1} - x_{t+1}}{R}\right) = R\,\beta\,E_t[V'_{t+1}(x_{t+1})]$$

Consider introducing a mean preserving spread into future incomes (i.e., more uncertainty but the same expected level of incomes). That is let future income be given by $\tilde{x}_{t+1} = x_{t+1} + \varepsilon_{t+1}$, where ε_{t+1} is a mean-zero random variable.

Applying Jensen's inequality we get:

$$R\,\beta\,E_{x\varepsilon,t}[V'_{t+1}(\tilde{x}_{t+1})] = R\,\beta\,E_{x\varepsilon,t}[V'_{t+1}(x_{t+1} + \varepsilon_{t+1})] > R\,\beta E_{x,t}V'(x_{t+1} + E_{\varepsilon}\varepsilon_{t+1})$$

$$= R\,\beta E_{x,t}V'(x_{t+1})$$

To maintain the equality in the f.o.c. the increase in future uncertainty forces a rise in cash on hand, x_t, via a fall in current consumption.

The logic for this increase in savings in the face of increased future risk is that prudent individuals suffer less pain from increased risk at higher levels of wealth. As a result, such individuals can decrease the pain from future risk by shifting wealth to the period with the risky income. In the precautionary savings model this implies shifting more wealth to the second period via an increase in saving.

Prudence can also be demonstrated by applying a second order Taylor approximation to the Euler equation for consumption:

$$u'(c_t) = R\,\beta\,E_t[u'(c_{t+1})]$$

$$u'(c_t) \approx R\,\beta\,E_t\left[u'(c_t) + u''(c_t)(c_{t+1} - c_t) + \frac{1}{2}u'''(c_t)(c_{t+1} - c_t)^2\right]$$

$$1 \approx R\,\beta\,E_t\left[1+\frac{c_t u'(c_t)}{u'(c_t)}\frac{c_{t+1}-c_t}{c_t}+\frac{1}{2}\frac{c_t u''(c_t)}{u'(c_t)}\ \frac{c_t u'''(c_t)}{u''(c_t)}\left(\frac{c_{t+1}-c_t}{c_t}\right)^2\right]$$

$\gamma(c_t)$=- $\frac{c_t u''(c_t)}{u'(c_t)}$ = Coefficient of relative risk aversion (RAA)

$\psi(c_t)$=- $\frac{c_t u'''(c_t)}{u''(c_t)}$= Coefficient of relative prudence (RP)

$$1 \approx R\,\beta\,E_t\left[1-\gamma(c_t)\frac{c_{t+1}-c_t}{c_t}+\frac{1}{2}\ \gamma(c_t)\psi(c_t)\left(\frac{c_{t+1}-c_t}{c_t}\right)^2\right]$$

$\gamma(c_t)$ and $\psi(c_t)$ are known at time t and since $c_{t+1} \approx c_t$, $E_t[c_{t+1}] \approx c_t$,:

$$E_t\left[\frac{c_{t+1}-c_t}{c_t}\right] \approx \frac{1}{\gamma_t}\frac{R\,\beta-1}{R\,\beta}+\frac{1}{2}\ \psi_t Var_t\left[\frac{c_{t+1}-c_t}{c_t}\right]$$

As $c_{t+1} \approx c_t$, $\frac{c_{t+1}-c_t}{c_t} \approx \Delta\log c_{t+1}$. Also, if we assume $R\,\beta \approx 1$ then $\frac{R\,\beta-1}{R\,\beta} \approx -\log R\,\beta$:

$$E_t[\Delta\log c_{t+1}] \approx -\frac{1}{\gamma_t}\log R\,\beta+\frac{1}{2}\ \psi_t Var_t[\Delta\log c_{t+1}]$$

The first term is the intertemporal motive for saving while the second is the precautionary motive. This latter term means a higher variance of consumption growth will result in higher expected consumption growth, lower current consumption and precautionary savings.

The precautionary motive can also be demonstrated in closed form using the Constant Absolute Risk Aversion (CARA) utility function. In the CARA utility function, $U(c_t) = -\exp^{-\gamma c_t}$, risk aversion is not related to the level of income of the consumer. Intertemporal substitution in this model takes the form of additive changes in growth of consumption instead of multiplicative changes in the CRRA model. This approach adds a variance term to the certainty equivalent model which represents prudence. The CARA Euler equation with uncertainty is:

$$\exp(-\gamma c_t) = \beta R E_t[\exp(-\gamma c_t)]$$

Assuming a linear policy rule for consumption, where A is the coefficient on consumption from assets, B is the coefficient on income and k is a constant:

$$c_t(a_t, y_t) = Aa_t + By_t + k$$

$$c_{t+1}\left(a_{t+1},y_{t+1}\right) = Aa_{t+1} + By_{t+1} + k$$

Substituting the evolution of assets ($a_{t+1} = R(a_t + y_t - c_t)$) and assuming that income follows a random walk, so that $y_{t+1} = y_t + \varepsilon_{t+1}$, gives;

$$c_{t+1} = [AR(1-A)]a_t + [AR(1-B)+B]y_t + (1-AR)k + B\varepsilon_{t+1}$$

Random walk income assumes that all income is permanent income in the certainty equivalent model and is consumed. Solve for *k* by substituting into the Euler equation:

$$e^{-\gamma c_t} = \beta R E_t\left[e^{-\gamma c_{t+1}}\right]$$

$$e^{-\gamma(Aa_t+By_t+k)} = \beta R E_t\left[e^{-\gamma([AR(1-A)]a_t+[AR(1-B)+B]y_t+(1-AR)k+B\varepsilon_{t+1})}\right]$$

First of all matching coefficients from both sides of the Euler equation gives:

A=AR(1-A) so A=r/R and;

B=AR(1-B) + B so B=1.

Then matching the constant:

$$e^{-\gamma k} = e^{-\gamma(1-AR)k}\beta R E_t\left[e^{-\gamma B\varepsilon_{t+1}}\right]$$

$$e^{-\gamma k} = e^{-\gamma\left[(1-AR)k-\frac{\log(\beta R)}{\gamma}-\frac{\log E_t\left[e^{-\gamma B\varepsilon_{t+1}}\right]}{\gamma}\right]}$$

Substituting A=$\frac{r}{R}$ and B = 1 (assuming an annuity rate is applied to assets and all income is consumed as it is permanent income):

$$k = -\frac{1}{r}\left[\frac{\log(\beta R)}{\gamma} - \frac{\log E_t\left[e^{-\gamma\varepsilon_{t+1}}\right]}{\gamma}\right]$$

Assuming $\varepsilon_{t|t-1}$is normally distributed with mean zero and variance σ^2:

$$E_t\left[e^{-\gamma\varepsilon_{t+1}}\right] = e^{\frac{\gamma^2\sigma^2}{2}}$$

$$k = -\frac{1}{r}\left[\frac{\log(\beta R)}{\gamma} + \frac{1}{2}\gamma\sigma^2\right]$$

Substituting for A, B and k gives the policy function:

$$c_t\left(a_t, y_t\right) = \frac{r}{R}a_t + y_t - \frac{1}{r}\left[\frac{\log(\beta R)}{\gamma} + \frac{1}{2}\gamma\sigma^2\right]$$

The first two terms are the same as for the certainty equivalent case; $\frac{r}{R}a_t$ is the annuity value of a_t, while all income y_t is consumed because it is permanent income from the random walk assumption. The third term reflects relative patience which decreases consumption. The difference with the certainty equivalent example is that here, increased variance or risk feeds into the additive variance term, which represents prudence.

2.4 Precautionary savings with borrowing constraints and the buffer stock model

An additional constraint is included in this problem to limit the resources available to the household to 'cash in hand' available at the start of the period (Deaton, 1991):

$$V_t(x_t) = \max_{\{c_t\}} u\,(c_t) + \beta\, E_t\,[V_{t+1}(x_{t+1})]$$

s.t.

$$x_{t+1} = R(x_t - c_t) + y_{t+1}$$

$$0 \leq c_t \leq x_t$$

The resulting Lagrangian now has an extra term due to the additional constraint:

$$V_t(x_t) = \max_{\{c_t\}} u\,(c_t) + \beta\, E_t\,[V_{t+1}(R(x_t - c_t) + y_{t+1})] + \lambda_t(x_t - c_t)$$

Where λ_t is the Lagrange multiplier associated with the constraint $c_t \leq x_t$. The f.o.c. is now;

$$u'(c_t) = R\,\beta\, E_t\,[V'_{t+1}(x_{t+1})] + \lambda_t$$

Using the envelope theorem, the marginal value of cash on hand is;

$$V'(x_t) = R\,\beta\, E_t\,[V'_{t+1}(x_{t+1})] + \lambda_t = u'(c_t)$$

It follows that when the borrowing constraint does not bind, the usual Euler holds;

$$u'(c_t) = R\,\beta\, E_t[u'(c_{t+1})]$$

When the borrowing constraint binds, $\lambda_t > 0, c_t = x_t$ and;

$$u'(x_t) > R\ \beta\ E_t[u'(c_{t+1})]$$

The borrowing constraint $c_t \leq x_t$ works in two ways:

- When it binds, the household is forced to consume less than desired:
- It also affects consumption now because it might bind in future. This is captured by $E_t[u'(c_{t+1})]$. The impact on consumption of fear of the borrowing constraint can be reinforced by the convexity of marginal utility and the precautionary savings motive. The more convex, the more likely the household will save now to reduce the likelihood of being constrained in future as insurance against income uncertainty, as the precautionary motive is stronger.

The buffer stock model adds the assumption that $\beta R < 1$, i.e. the household is impatient and would choose to consume more today and run down assets in the absence of the precautionary motive. The result is that the household will aim to achieve a certain target level of liquid assets and not more. Once households have achieved this target level of wealth, consumption will track income, which is consistent with empirical research finding 'excess sensitivity' of consumption to income. The dynamics behind the buffer stock model are illustrated in the phase diagram given in Figure 2.5 below, which is taken from Carroll and Toche (2019). The Δa^e=0 locus indicates for a given level of wealth a^e, how much consumption c^e would leave a^e unchanged. Consuming less than the locus level for any given a^e will cause wealth to rise, and thus the horizonal arrows point left above the consumption locus and right below. The precautionary motive affects the Δc^e=0 locus, where growth in consumption is zero. In the Carroll and Toche presentation, where uncertainty only relates to unemployment, c^e is consumption while employed and c^u is consumption while unemployed. The Euler equation determines the amount by which consumption will fall during unemployment. However, similar to the borrowing constraints example where the constraint binds, increased wealth increases the level of consumption during unemployment, and thus reduces c^e_{t+1}/c^u_{t+1}. Thus where wealth is increasing to the right of the locus, c^e_t is declining (hence the down arrow in the phase diagram) and vice-versa.

Figure 2.5 Illustration of the buffer stock model

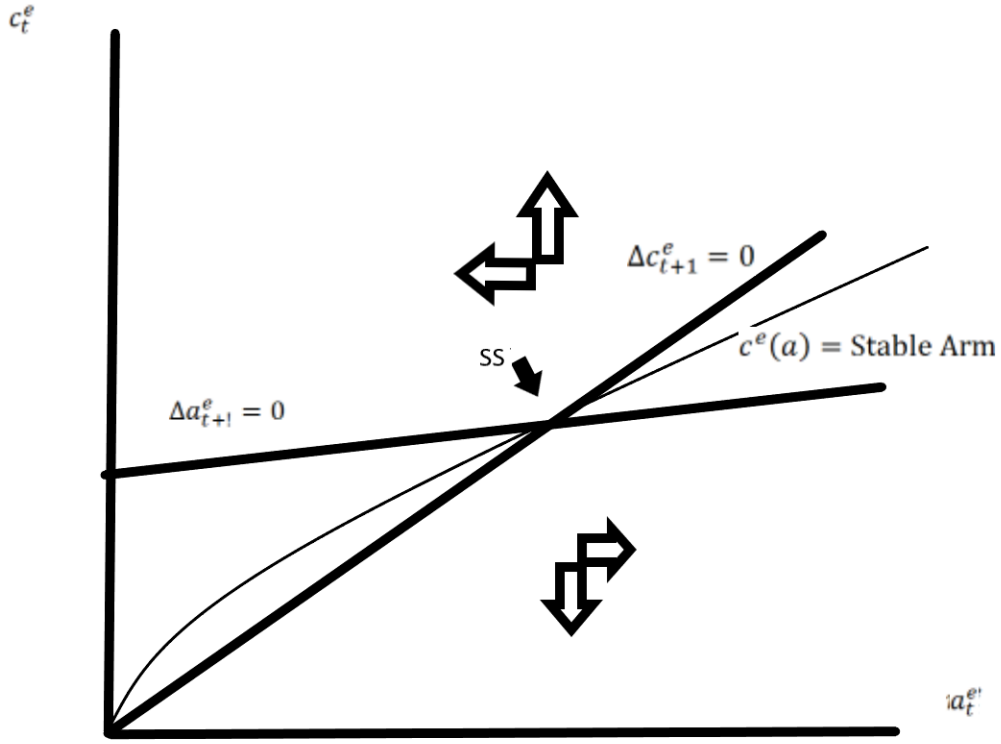

In summary, the main characteristics of life cycle models with precautionary savings, borrowing constraints and buffer stocks are;

- At any given level of income consumption is lower in these models since uncertainty results in precautionary savings. Consumption growth will be higher over time as the consumption profile is tilted up by precautionary savings. Consumers with lower levels of assets or income react more strongly to the precautionary motive, as they are more exposed to the risks from bad outcomes.
- The differences between the three models are clearest for these consumers with low assets or incomes. They are less inclined to smooth their incomes by borrowing, which means that consumption is more likely to track income. This is most clearly demonstrated where the borrowing constraint binds. The ratio of patience to the interest rate is generally less relevant in this case, and for low wealth consumers it has very little impact on decision making. The buffer stock model in particular requires that consumers are impatient. Under certainty equivalence, these consumers would go into

debt. In the buffer stock model, everything above the buffer level of savings is consumed at lower income levels.

- Consumers have a higher marginal propensity to consume than in the certainty equivalent models as the precautionary motive weakens as income or wealth rises. This is particularly relevant for transitory incomes;
- Consumers are either reluctant to borrow or cannot borrow so consumption levels will be affected by predicted and unpredicted shocks. They are much less 'forward looking' which exacerbates the impact of shocks, since increased uncertainty directly affects current consumption. Consumption is strongly related to predictable earnings growth in these models (Carroll, 1992);
- The precautionary motive and fear of borrowing constraints also affects retirement consumption, as pensioners will have a fear of medical expenses and uncertainty about remaining life years. Buffer stock models assume that pensioners maintain or increase their buffer level of savings in retirement. These models also have implications for retirement planning, as consumers are unlikely to start saving for retirement at young ages if they are impatient.

The purpose of this chapter is to illustrate the main theoretical consumption models used in the literature. My subsequent chapters focus on deviations from the 'certainty equivalent' model due to measurement issues, the age profile of consumption and formation of expectations. I also examine the impact of income volatility on consumers and how households insured themselves against this volatility to maintain their living standards during the "Great Recession".

3. Income, expenditure and consumption in the Household Budget Survey

3.1 Introduction

International studies have raised concerns about measurement error for incomes in social surveys, especially at low income levels. Meyer and Sullivan (2003, 2009, 2013) have highlighted this issue in the US in the context of difficulties with measuring means tested and non-cash benefits for official poverty measures. Brewer, Etheridge and O'Dea (2017) have raised the same issue with UK income data, drawing attention to unrealistically low incomes compared to reported expenditure for the same households. Concerns have also been raised about underreporting of expenditure in household surveys based on National Accounts comparisons, with Garner et al. (2006) analysing US data in detail. Studying the relationship between income and expenditure requires accurate measures of both. This paper considers these issues in the Irish context.

This paper examines (i) measurement error in reported incomes in the Irish Household Budget Survey (HBS) and (ii) the distribution of consumption based on new estimates using HBS expenditure data. The HBS is a sample survey of households in Ireland, which captures income and expenditure for the same households. My results show that income appears to be better measured at low levels in the HBS than in similar UK/US surveys. A change in the HBS survey process in 2009/10 facilitates an experiment which suggests that better measurement is due to the relatively simple Irish benefits system. This analysis also highlights more consumption smoothing in 2009/10 than in 2004/5, and that an unemployed group who were 'transitory income poor' replaced a lone parent group who were 'permanent income poor' at the bottom of the income distribution between these years.

This paper is structured as follows. Section 3.2 provides a literature review. Section 3.3 describes the data used for the analysis and Section 3.4 sets out the methodologies applied from the literature. Section 3.5 summarises the research findings, with conclusions in Section 3.6.

3.2 Literature review

As discussed in detail in Chapter 2, the 'Permanent Income/Life-Cycle Hypothesis' proposes that Permanent Income or consumption (which depends only on Permanent Income) yields a

more appropriate measure of the distribution of household resources than current income. Permanent Income is the annuity value of expected lifetime wealth based on information currently available to consumers. Under the Permanent Income theory, households smooth their lifetime consumption in response to short term income shocks or 'transitory income' shifts. Insurance against shocks is usually facilitated through running down savings or borrowing. Over the longer term, households maintain 'life-cycle' consumption through a similar savings process. There have been various strands to economic research of interest for my analysis, which can be broadly grouped into two categories for the purposes of this paper; research on population experience related to the 'Permanent Income/Life-Cycle Hypothesis' and research on measurement issues with income, expenditure and consumption.

Empirical tests of the Permanent Income and Life Cycle Hypotheses suggest that where households can 'consumption smooth' then consumption is a better measure of resources than current income. Economists try to identify evidence of insurance and savings in household data when testing the theory. Until Hall (1978), much of the research on the theory was based on macroeconomic analysis of the relationship between lagged income and consumption. Hall proposed that consumption growth reflects expected future income, and thus only responds to unexpected events and is unpredictable. He found empirical evidence for his proposal from the US National Accounts. Hall's proposal provided a theoretical basis for using repeated cross section data to test the theory, as the relationship between current income and current consumption is sufficient for empirical tests. Deaton and Paxson (1994) tested Hall's conclusion that the consumption of each person in the economy follows a random walk. They found that the variance of earnings, income and consumption increased over time within cohorts in keeping with the Permanent Income Hypothesis using repeated cross section 'pseudo panel' data for the US, UK and Taiwan. Blundell and Preston (1998) found that this increase in permanent variance will result in higher levels of aggregate inequality over time.

Leaving theoretical considerations aside, measurement issues have become much more prominent in the income vs. consumption debate in recent years. Meyer and Sullivan (2003) found that expenditures significantly exceed incomes up to the 20th percentile of each distribution. Based on data from the Consumer Expenditure Survey (CES – the US equivalent to the HBS/FES) reported by the same individuals, average expenditures in the bottom decile of incomes were over 4.6 times average income. The CES was also found to significantly under report incomes compared to official levels of means tested benefits (Temporary Assistance for Needy Families and Food Stamps). Brewer, Etheridge and O'Dea (2017) examined income

measurement in the UK using a similar set of empirical strategies to Meyer and Sullivan (2003, 2009) and also found evidence of underreporting of incomes. They additionally produced modelled estimates of income given expenditure and vice-versa. They found that expenditure levels for the very lowest incomes in the UK survey were implausible as there is a distinct 'tick' in spending given income in the bottom 1 or 2 percentiles of incomes. Figure 3.1 below is copied with permission.

Figure 3.1 Median expenditure by income and median income by expenditure, eq. £/wk

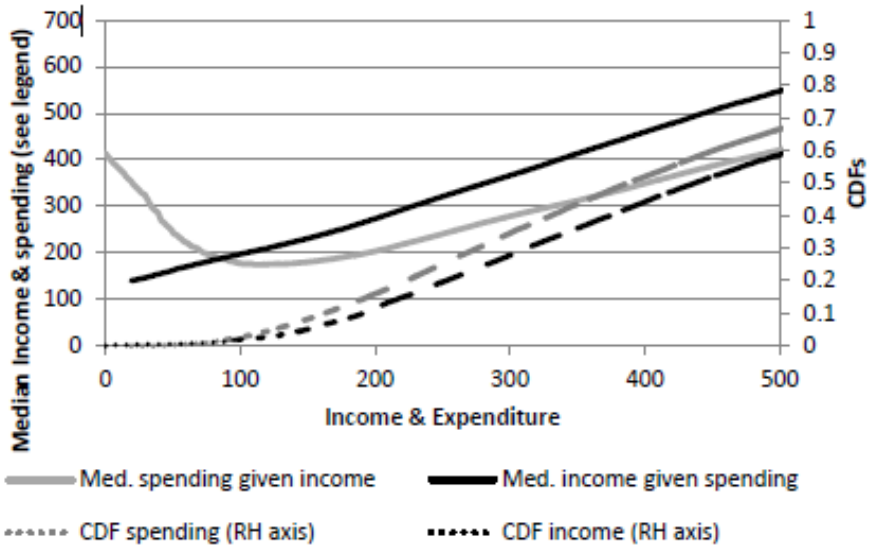

Notes: Graph shows median spending given income, and median income given spending, both drawn using a locally-weighted median regression. Data: Living Costs and Food Survey 2006/07-2009 as described in Section 3.

Spending levels in the bottom percentile of equivalised income are estimated at £400 per week which is more in line with the 60th percentile of income than the first. Undercoverage of tax credits in reported survey incomes by respondents is presented as the main reason for mismeasurement of low incomes. They compare survey estimates of tax credits to official costs and find that just half to two-thirds of the official income is covered in the survey. Using panel data, they also examined transitions in and out of low incomes, for evidence of consumption smoothing of transitory shocks, and found that the tick could not be fully explained by the income transitions or by economic models of saving and borrowing. However, unlike Meyer and Sullivan (2003, 2009, 2013), they stop short of recommending expenditure or consumption over income for poverty measurement due to concerns about underreporting of expenditure in household surveys when compared with National Accounts consumption data.

Opinion has changed in the UK literature about the nature of undercoverage of consumption in FES expenditure compared to National Accounts consumption. Tanner (1998) found that undercoverage was largely due to methodological differences between the two sets of statistics, and expenditure was accurately reported in the FES for categories that were measured on the same basis in the National Accounts. By contrast, Blundell and Etheridge (2010) found that the FES measured income well while expenditure was increasingly underreported compared to National Accounts Personal Consumption Expenditure (PCE). This suggests measurement error for expenditure in the survey is increasing over time.

Garner et al. (2006) also examined this issue in detail for the US and found three main reasons for undercoverage of PCE in the CES; scope, definition and methodology differences. PCE includes expenditures of non-profit institutions, which extends both the scope and definition of expenditures recorded in the CES. For methodology, the CES is collected via household surveys while PCE estimates are mainly based on business surveys and administrative data. They found that the ratio of CES to PCE expenditures fell from 0.88 in 1992 to 0.82 in 2003 after taking account of scope, definition and methodology differences. Aguiar and Bils (2015) related these concerns back to theory-based explanations of the gap between consumption and income inequality growth, which originated with Deaton and Paxson. They found that when systematic undercoverage in CES expenditures is 'corrected' with PCE weights allowing for scope, definition and methodology differences, the discrepancy in growth between mean income and mean expenditure is reduced by 40%.

How does the choice of income or consumption identify risk groups in the population in practice? Poterba (1991) presented an analysis of the distribution of gasoline taxes in the US based on insights from the life-cycle and permanent income theories of consumption using 1985 CES data. Expenditure on gasoline as a percentage of income displayed a regressive pattern, with higher shares for low income groups. Poterba found the opposite pattern for gasoline as a percentage of expenditure when households were ranked by expenditure rather than income, which he related to smoothing of income fluctuations in permanent income. Further analysis found that the elderly and families with children were likely to spend more than they earned, as were groups experiencing hardship, all of which is consistent with the permanent income model. Cutler and Katz (1992) broadened the Poterba analysis to examine changes in the distribution of consumption and income in the US over the 1980s. They found

that changes in the distribution of consumption[4] correspond closely to changes in the distribution of income. Consumption is more equally distributed than income in each year, in line with the permanent income hypothesis, but both measures became more unequal during the 1980's. In addition, the risk groups identified by both measures are similar. They concluded that there is no 'hidden prosperity of the poor' in the consumption data. Slesnick (1993) followed a similar methodology to Cutler and Katz over the same period but focussed on poverty rates for income and consumption rather than inequality. His key finding was that poverty rates were substantially lower throughout the period for consumption than income regardless of the equivalence scales or poverty thresholds used. This is consistent with the Permanent Income Hypothesis and suggests that the level of resources needed for the 'war on poverty' was lower than as argued by liberal commentators at the time. The consumption-based poverty group had characteristics in line with those of the 'permanent income poor', as they had fewer physical assets, high shares of necessities in family budgets and low levels of dissaving.

O'Neill and Sweetman (2001) compared a range of poverty, distribution and inequality measures for expenditure and income in Ireland using data from the HBS in 1987 and 1994. They found that the distributions of income and expenditure were remarkably consistent in both years, but that the composition of the group identified as poor was different for both measures, in line with the life-cycle/permanent income model. In particular, farmers and the self-employed fared better with expenditure than income and are known to face a variable and uncertain income stream. Krueger and Perri (2006) updated the earlier US analysis of CES data to the 1980-2003 period and found that income inequality strongly increased in the 1990s while consumption inequality was flat. They found that the difference was due to a divergence in within-group inequality over the period, which is in line with Deaton and Paxson's analysis. Within-group inequality includes inequality caused by income shocks, which can be 'smoothed' in consumption. Between-group inequality arises from changes in returns to observable characteristics of the household, such as education levels. It is unlikely that households can insure against changes in these returns so between-group income and consumption inequality should, and did, track each other.

[4] Using expenditure from the CES with adjustments in line with economic theory similar to the analysis in this paper.

Meyer and Sullivan (2009) found significant differences in poverty rates and risk groups using a range of resource measures derived from US income and expenditure data. While saving and dissaving contributes to differences in outcomes between income and consumption over time, they presented detailed evidence similar to their 2003 paper identifying increasing mismeasurement of incomes as the main factor behind the differences. Brewer and O'Dea (2012) found similar results based on UK data, and both studies raise concerns with the reliability of income statistics for official poverty measurement in the US and UK. Blundell and Etheridge (2010) also examined repeated cross sections of UK incomes data from 1978-2005, finding similar results to the Blundell and Preston (1998) 'inequality boom' in the 1980s. However, they cautioned that FES expenditure increasingly understated consumption from the National Accounts in the UK, as mentioned earlier. They extended the Blundell and Preston (1998) analysis to include earnings, employment, taxes and transfers in an attempt to link microeconomic and macroeconomic factors behind the evolution of inequality. This analysis identified a corresponding 'inequality boom' during the 1980's due to changes in skill prices. Meyer and Sullivan (2013) extended their earlier US poverty analysis to include inequality measurement, and found, similar to Kreuger and Perri (2006), that income inequality increased much more strongly than consumption inequality since the 1980's. Their additional contribution to the research was to show that the divergence is concentrated in the bottom half of the two distributions.

3.3 Data

The anonymised HBS datasets lodged with the Irish Social Sciences Data Archive (ISSDA) for the survey years 1999/2000, 2004/5 and 2009/10 are used for this analysis. The income information collected in the survey is secondary to the main purpose of the HBS, which is measurement of household expenditure in order to construct price index weights. The HBS includes detailed information on income receipts by source, expenditure patterns and household composition for nationally representative samples of Irish households. While incomes are of secondary importance to expenditure data, the income distribution reported from HBS is very similar to the Survey of Income and Living Conditions (SILC), which is the main survey used for income measurement in Ireland[5]. This indicates that HBS is suitable for incomes analysis. For incomes, household disposable incomes are used in the analysis in this

[5] https://www.cso.ie/en/media/csoie/methods/surveyonincomeandlivingconditions/A_Consistency_Check_between_SILC_2010_and_HBS_2010.pdf

paper which take account of taxes and social welfare incomes; total reported expenditure excluding repayments of personal loans other than mortgages is used for expenditure; and consumption is constructed in this paper from HBS expenditure. Expenditure in HBS is based on a two-week diary of household purchases with a 'lookback' for large occasional purchases over the previous year. Household incomes, expenditure and consumption are equivalised in this paper using the National equivalence scale from ESRI/CSO analysis of SILC incomes, which assigns a weight of 1 for the first adult, 0.66 for the second adult and 0.33 for children aged 14 and under[6]. The HBS sample sizes for 1999/2000, 2004/5 and 2009/10 were 7,644, 6,870 and 5,891 households respectively.

In addition to the expenditure and incomes data already available in the HBS, two supplementary sources (National Accounts and Census of Population) are used in this paper to replicate the Cutler and Katz (1992)/Meyer and Sullivan (2009)/Brewer and O'Dea (2012) analyses of consumption. Both MS and BoD compared expenditure from the US/UK versions of HBS with PCE estimates from the National Accounts at commodity level, with a focus on accurate representation of consumption at the lower end of the distribution. Garner et al. (2006) gave a detailed analysis of the measurement and methodological differences between CES expenditures and PCE, which are also relevant in Ireland. PCE is based on a fully exhaustive compilation process, using a mix of administrative and survey data. Most of the data sources for PCE are business surveys and Revenue register data on trade, while HBS is a sample household survey. The key data source for PCE is 'commodity flow', i.e.:

Consumption = Production + Imports + Mark Up – Exports.

While HBS data are used in some of the detailed PCE calculations, generally PCE is independently estimated. Also, PCE is measured according to economic concepts of consumption, including service flows from assets and non-market expenditures for households' benefit, while expenditure simply reflects cash outlays. Garner et al. (2006) gave a detailed presentation of the scope, definition and methodology differences between the two measures.

Microdata from the Census of Population are used in this analysis to calculate market rent estimates for a grid of housing types. The information collected in the Census includes rent,

[6] There are a large number of possible equivalence scales; Callan et al. (1996) argue that analysis of poverty over time is not sensitive to the choice of scale. They also say that there is a strong case for using the person as the unit of analysis rather than the household, mainly for accurate presentation of incidence by population sub-groups – that is not in scope for this paper and thus equivalised household income/expenditure is the unit used. More recent research (Regan and Kakoulidou, 2022) suggests that poverty trends, particularly at disaggregated levels, are somewhat sensitive to the choice of scale.

size of household and house type, among other details. Rents are used as the economic value of living in each particular house type – i.e., the 'service flow' from living in a house instead of mortgage payments or purchase costs. Most of the relevant US and UK research used administrative sources to estimate these service flows, but the approach here is similar to that used in the compilation of the Irish National Accounts.

3.4 Methodology

The methodology for the paper is based on the Irish, UK and US studies described in the literature review. Three specific methodologies from the literature are applied to the Irish data;

1. Calculation of consumption estimates using HBS expenditure data (Cutler and Katz, 1992, Meyer and Sullivan, 2009 and Brewer and O'Dea, 2012);
2. Locally weighted median regression estimates of income given expenditure and expenditure given income (Brewer, Etheridge and O'Dea, 2017);
3. Kernel density estimates (O'Neill and Sweetman, 2001).

3.4.1 Consumption estimates based on HBS expenditure

HBS expenditure is compared with PCE consumption for the same products to assess the accuracy of expenditure and also as a guide for constructing a consumption measure. Table 3.1 replicates the MS/BoD comparison of HBS and PCE using Irish HBS expenditure adjusted to calendar year values and annual PCE data converted to weekly amounts per household. Overall, there is no evidence of growing undercoverage of consumption by HBS as found in the US by Garner et al. (2006) and in the UK by Blundell and Etheridge (2010). When product groups with methodological or measurement issues are excluded in the last row of the table, the differences between expenditure and PCE consumption are quite low. This suggests specific rather than systematic undercoverage in the HBS data. Each of these specific groups are discussed below.

Table 3.1 Weekly HBS expenditures compared to PCE by product group and year

Group code	Product group	HBS 2000 €/Weekly	PCE 2000 €/Weekly	HBS 2005 €/Weekly	PCE 2005 €/Weekly	HBS 2010 €/Weekly	PCE 2010 €/Weekly
1	Food (including meals out)	€117.67	€121.36	€142.74	€137.51	€131.28	€133.41
2	Alcohol and tobacco (including out)	€44.10	€125.07	€47.18	€139.18	€39.48	€126.14
3	Clothing and Footwear	€35.11	€49.03	€42.67	€46.24	€40.11	€35.11
4	Housing and utilities	€77.09	€136.66	€125.16	€194.25	€183.08	€199.83
5	Furnishings and household equipment	€41.27	€49.36	€52.97	€54.94	€46.55	€35.54
6	Health	€19.65	€31.16	€33.93	€56.39	€37.34	€71.21
7	Transport including vehicles	€94.92	€101.33	€122.74	€125.89	€116.31	€116.32
8 to 13	Miscellaneous services	€147.91	€190.27	€219.68	€265.17	€216.47	€247.32
Total	Total Expenditure/Consumption	€577.72	€804.22	€787.07	€1,019.57	€810.62	€964.88
	Of which; for comparable groups	€288.97	€321.07	€361.12	€364.57	€334.25	€320.38

The main adjustments from the literature applied to expenditure to calculate consumption are replacement of cash expenditures with service flows and inclusion of non-profit expenditures. As with MS/BoD for the US and UK, the adjustments for these items in the consumption estimates calculated in this paper are particularly concentrated in lower income deciles. This is mainly due to policy design for income supports such as medical services and social housing. Free medical services are provided to low income households via 'medical cards' in Ireland. Rents for social housing are also subsidised for low income groups. However, life-cycle income patterns also play a role – pensioners tend to have low incomes but high levels of imputed rent. The first and most significant of the MS/BoD adjustments is imputed rent amounts for owner occupiers instead of reported mortgage payments in Group 4. Vehicle service flows are once again used instead of reported vehicle payments in Group 7, also based on the economic concept of service flows from assets. An imputed value is included in Group 6 consumption for free medical services provided by Government. In addition to these methodological adjustments, an adjustment for underreporting of alcohol and tobacco is included. This differs from MS/BoD as they only adjust for methodological differences between the two measures. Undercoverage of these products is less of an issue in the CES and FES (e.g., Tanner reported 60-66 per cent coverage in the FES). In the HBS, on the other hand, the difference is significant and clearly needs to be addressed. The adjustments are described in detail below.

Finally, there are many smaller methodological issues which result in undercoverage of 'Miscellaneous Goods and Services' (Groups 8-13) using HBS expenditure as a consumption estimate. These items include the adjustment for investment income as insurance consumption

for households, State support for student fees and various other Social Transfers in Kind (STiKs). These differences are either not large enough individually or widely distributed enough to consider as specific adjustments in this paper.

A. Adjustment for housing

Imputed rent attributes an equal consumption value to households of a given housing quality for all tenure types. The adjustment for imputed rent used in the consumption measure is a simplified version of the calculation used in the Irish National Accounts. Private rents for tenants from the 2002, 2006 and 2011 Census of Population (CoP) are used as the basis for this calculation. CoP is a comprehensive source of rent data for the population and includes details of housing quality for stratification of the reported rents. CoP rents replace mortgage expenditures for owner occupier households and rents for social tenants in the HBS microdata for 64 house types stratified by location, number of rooms and house type.

Table 3.2 Imputed weekly rent calculations based on 2011 Census of Population

Rooms and location classification		Detached House	Other House	Purpose Built Apartment	Other Apartment
		€/Weekly	€/Weekly	€/Weekly	€/Weekly
1-3 rooms	Small urban	€98	€99	€101	€78
	Large urban	€105	€106	€109	€88
	Dublin	€152	€153	€166	€118
	Rural	€84	€88	€93	€80
4 rooms	Small urban	€116	€112	€122	€99
	Large urban	€123	€119	€124	€110
	Dublin	€191	€187	€198	€187
	Rural	€103	€103	€112	€97
5 rooms	Small urban	€122	€120	€138	€113
	Large urban	€130	€127	€148	€138
	Dublin	€202	€198	€232	€222
	Rural	€107	€107	€125	€118
6+ rooms	Small urban	€152	€136	€142	€103
	Large urban	€175	€160	€151	€126
	Dublin	€283	€245	€208	€150
	Rural	€146	€129	€145	€119

The Consumer Price Index rent index is used to interpolate inter-censal years to match the survey reference periods for HBS. The effect of the adjustment for imputed rents (i.e, the difference between mortgage payments and the rents that replace them in the consumption measure) is higher in 2004/5 than in 2009/2010. This reflects the rapid increase in HBS mortgage expenditures between the two years, and also lower private rents. Estimates for 2009/10 based on the 2011 CoP are shown in Table 3.2. The cell structure used for the CoP data was chosen to match the level of detail available for households in the HBS microdata from ISSDA.

B. Adjustment for medical services

An imputed value of in cash terms of €39 in 2009/10, €24 in 2004/5 and €21 in 1999/2000 per household is applied to private health insurance costs for households with medical cards who are entitled to free services. This value equalises total medical expenditure, excluding 'out of pocket' costs, between households with and without medical cards. There is evidence that medical services consumption is higher for households with medical cards (c.f. Madden, Nolan and Nolan, 2005) but consumption is assumed to be similar to non-medical card holders as it is not obvious that living standards are higher for people who are more likely to have medical conditions. There are difficulties in assigning monetary values to non-cash benefits (O'Dea and Preston, 1994) as the value differs by household depending on usage and also the value each household attributes (e.g., 'willingness to pay') to the service provided.

C. Adjustment for cars

The economic rationale for this adjustment is similar to that used for imputed rents. While there is no significant difference between HBS and PCE for Transport including vehicles (Group 7), the literature proposing consumption-based resource measures (MS/BoD/Cutler and Katz) indicates that expenditure is a poor basis for measuring the consumption value of large asset purchases such as cars from which an economic flow value is derived over years by consumers. Table 3.3 shows the annual consumption value per car per household that is used instead of expenditure to represent a flow consumption value**.** Most households with cars report 'zero cost' in the HBS, particularly in low spending households. I use the variation in car expenditure between expenditure deciles as a quality adjustment factor. Weekly expenditure on cars in each HBS expenditure decile is divided by the number of cars owned in that expenditure decile to reallocate the expenditure across all car owners. As an example, in the first decile of expenditure in Table 3.3, 86% of the 44% of households with cars are assigned a replacement

consumption estimate of €14.51 per car owned by the household rather than the zero expenditure reported in the HBS, as are the remaining 14% who record positive expenditures.

Table 3.3 Basis for cars consumption estimate 2009-2010 by expenditure decile

Expenditure decile	% of households with cars %	Total cars in decile N	% of households with zero costs %	Imputed expenditure per car €
Decile 1	44%	293	86%	€14.51
Decile 2	67%	431	83%	€9.87
Decile 3	76%	517	79%	€13.06
Decile 4	83%	590	74%	€15.46
Decile 5	87%	660	70%	€19.40
Decile 6	88%	731	61%	€26.42
Decile 7	90%	786	62%	€27.58
Decile 8	91%	851	61%	€31.97
Decile 9	94%	919	56%	€41.07
Decile 10	96%	1,108	56%	€39.07

D. Adjustment for alcohol and tobacco

The adjustment for understatement of alcohol and tobacco expenditure compared to PCE estimates is applied pro rata to households that report such expenditures in HBS. In other words, consumers are assumed to uniformly underreport expenditures rather than reporting missing values. The adjustment factor is based on the difference between PCE and HBS Product Group 2 expenditures in each survey year. The logic here is not fully consistent with the adjustment for medical services, where additional consumption was not considered for people likely to have below average health. Although alcohol and tobacco consumption are likely to result in negative welfare spillovers, the purpose of this adjustment is to account for HBS measurement error.

These four adjustments allow me to construct a measure of consumption from HBS expenditure which I use in later analysis.

3.4.2 Locally weighted median regression and Quantile Regression

I begin by replicating the Brewer et al. (2017) analysis using Irish data. They estimate a locally weighted median regression of expenditure given income, and vice versa to test for mismeasurement of expenditure. As with the UK analysis, this will assess whether low incomes are realistic in the context of reported expenditure.

The locally weighted or nonparametric median regression technique produces a series of regression estimates, each of which uses a subset of the observations that are close to a set of target points. More weight is given to observations that are close to the target points. In the context of expenditure given income, the locally weighted median regression gives a regression based point estimate of expenditure for each level of actual income chosen.

Quantile regression was introduced for modern econometric use by Koenker and Bassett (1978). The approach is based on sorting and ordering observations in a population. Calculation of an unconditional population quantile can be written as an optimisation problem to illustrate the approach. For any 0< τ<1, a linear function $\rho_\tau(u) = u(\tau - I(u < 0))$ is defined to set up the optimisation problem:

$$R(\varepsilon) = \min_{\varepsilon \, \epsilon \, R} \sum_{i=1}^{n} \rho_\tau(y_i - \varepsilon) \qquad (1)$$

This is the minimised sum of asymmetrically weighted absolute residuals with weights τ above quantile level τ and $\tau - 1$ below. This problem returns the τ^{th} population quantile as its solution. This can be observed through the directional derivatives with respect to ε over a distance *h*;

$$R'(\varepsilon+) = \lim_{h \to 0+} \frac{\big(R(\varepsilon + h) - R(\varepsilon)\big)}{h}$$

$$= \sum_{i=1}^{n} \lim_{h \to 0+} \frac{\big(\rho_\tau(y_i - \varepsilon - h) - \rho_\tau(y_i - \varepsilon)\big)}{h}$$

Substituting $y_i - \varepsilon - h$ into the linear loss function for ρ_τ;

$$\rho_\tau(y_i - \varepsilon - h) = (y_i - \varepsilon - h\,)(\tau - I(y_i - \varepsilon - h < 0))$$

Approaching zero from above the condition $y_i - \varepsilon - h < 0$ equals $y_i - \varepsilon < 0$;

$$\rho_\tau(y_i - \varepsilon - h) = (y_i - \varepsilon - h\,)(\tau - I(y_i - \varepsilon < 0))$$

Giving a numerator for the directional derivative of;

$$(y_i - \varepsilon - h\,)(\tau - I(y_i - \varepsilon < 0)) - (y_i - \varepsilon)\,(\tau - I(y_i - \varepsilon < 0))$$

Giving a complete expression for the directional derivative calculation;

$$\mathrm{h}((I(y_i \leq \varepsilon) - \tau)/\mathrm{h}$$

$$= \sum_{i=1}^{n} (I(y_i \leq \varepsilon) - \tau) \qquad (2)$$

And similarly for the left derivative;

$$R'(\varepsilon -) = \sum_{i=1}^{n} \big(\tau - I(y_i \leq \varepsilon)\big) \qquad (3)$$

None of the directional derivatives can be negative for an optimal solution since the two vector gradients are no longer converging towards the quantile level τ, and thus τ is optimal. To obtain an estimate of the conditional quantile function, I replace the absolute values in ε in the Equation (1) by the parametric function $\varepsilon(x_i, \beta)$, which can then be solved by linear programming methods;

$$Qy(\tau|x) = \frac{\min}{\beta \in Rp} \sum_{i=1}^{n} \rho_\tau \left(y_i - \varepsilon(x_i, \beta)\right)$$

Median regression is a special case with τ=0.5 and symmetrical absolute residuals.

Quantile regression assumes a linear parametric form for the data generating process underlying the conditional quantile estimates, i.e. $\varepsilon(x_i, \beta)$ is a linear function with a fixed form for covariates. However, in this chapter a nonparametric median regression is implemented for the main analysis to let the data drive the functional form. This approach uses local weighting in order to accurately estimate expenditure given income at each point of the income distribution. Weighted parametric (median) regressions are run at each selected 'target' point of income or expenditure to produce estimates of the conditional median expenditure or income at each point. Further weighted parametric regressions are run at increasing target values for an estimate of the non-parametric function. There are two choices required to perform this analysis;

(i) Bandwidth – the range of observations used to calculate the regression estimates. A window of 25% of observations was chosen, i.e. 25% of all observations around the target point were given weight in the estimates;

(ii) The kernel function – this assigns a weight of 0 to 1 based on the distance from the target point at the centre of the 25% window of observations. The tri-cube weight function was used in this analysis;

$$w(x) = (1-|d|^3)^3$$

Where *d* is the distance of a given data point from the point on the curve being fitted, scaled to lie in the range from 0 to 1.

A median regression of expenditure on income based on the set of income data points generated by the kernel function was then implemented, using the regression approach described earlier. Weighted estimates of equivalised expenditure were produced based on locally weighted equivalised incomes in €10 increments of target income from €50 to €150 and then in increments of €100 from €200-800. The same approach was followed for income given expenditure.

3.4.3 Kernel density estimates

Kernel density estimates have been widely used in economic research as a means of visually presenting summary information on distributions. Rosenblatt (1956) and Parzen (1962) independently developed the technique. Kernel density estimators approximate the density *f(x)* from observations on *x:*

$$f(x) = \frac{1}{nh}\sum_{i=1}^{n} K\left(\frac{x - X_i}{h}\right)$$

Where h is the window width or bandwidth and K is the kernel estimator, both described in Section 3.4.2. The presentation is similar in effect to histograms, which report the number of data points between non-overlapping intervals in 'bins.' For kernel density estimates, the range is still divided into intervals but the intervals are allowed to overlap according to the 'window' described earlier. This facilitates a much richer presentation of frequency data, as the summary of the distribution is smooth and independent of the choice of origin. This allows for standardised comparisons of distributions of different variables.

Bandwidth and kernel function choices are also needed for this approach. The Epanechnikov function with a chosen bandwidth of 50 was chosen as it presented the data with a good balance between smoothness and showing the peaks in each distribution.

3.5 Data analysis

3.5.1 Measurement error in the HBS data

I begin my analysis by considering the issue of measurement error in income. To do this I follow the procedure used in Brewer et al. (2017). This approach involves running a locally

weighted median regression of expenditure on income and examining how expenditure changes as income changes. Figure 3.2 presents 2009/10 HBS estimates for Ireland for comparison against the UK findings.

The main finding is that the Irish results show a much smoother progression than found in the UK for expenditure given income. This means that expenditure is more consistent with low incomes for a given household than in the UK analysis. There is no evidence of mismeasurement of incomes in HBS in Figure 3.2. However, as the UK 'tick' was found for very low income levels, more detailed analysis at the bottom of the income distribution is needed for a true comparison of the findings.

Figure 3.2 Median expenditure by income, and median income by expenditure, HBS 2009/10

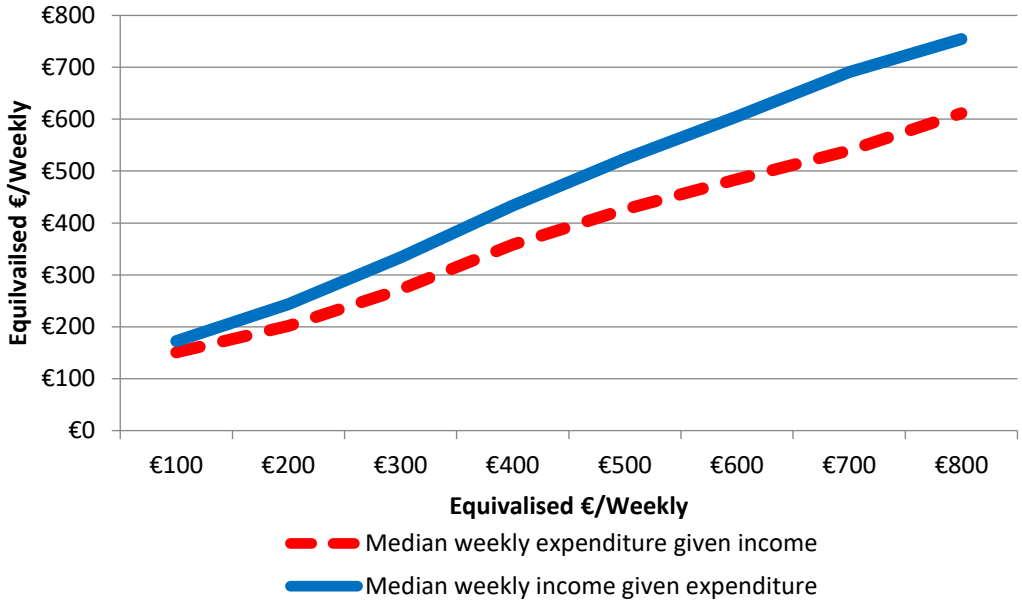

Figure 3.3 presents more detailed estimates of modelled expenditure from the bottom decile of equivalised income. There is still no 'tick' in expenditure given income at these extremely low income levels, while there was clear evidence of mismeasurement at this level of detail in the UK data. The cumulative distribution functions for income and expenditure are in Figure 3.11 in the Appendix, which shows a higher share of expenditure than incomes at each of the levels reported with both distributions converging at 90-92 per cent at €800. Income is more right skewed than expenditure on this basis.

From 2009/10, social welfare incomes were collected directly from administrative sources in the HBS. This process change could be expected to improve the accuracy of the low income

results and may help to explain the difference with UK findings. 2004/5 HBS welfare data were collected via survey methods, which is more in line with the UK approach. 2004/5 expenditure given income is also included in Figure 3.3, as it facilitates a comparison of results before and after the data collection change that is additional to the international research. If the change in data collection leads to apparent better measurement, then there should be a 'tick' in the 2004/5 results. However, while the shape of expenditure given income is different in 2004/5 and 2009/10, there is still clearly no 'tick' for the lowest incomes. This analysis suggests that the method of data collection is not responsible for the reliability of low incomes data in the HBS. Complexity is the US and UK benefits systems is raised as the main reason for mismeasurement of low incomes in the literature, and the Irish benefits system is much simpler as it is based on flat rate cash payments.

Figure 3.3 Median expenditure by income, HBS 2009/10 and 2004/5

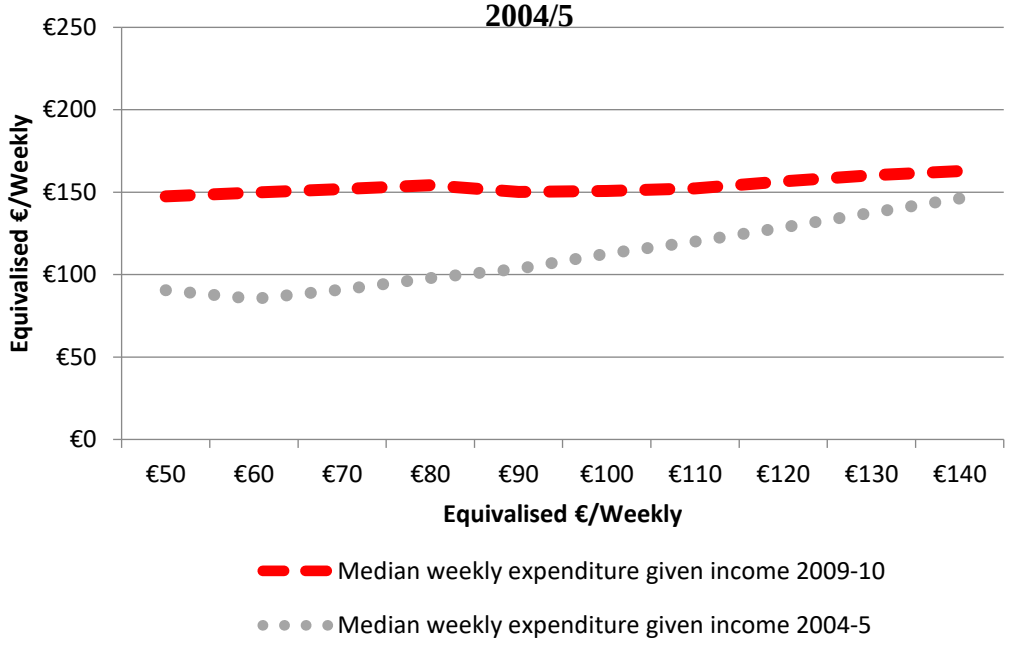

The analysis in Figure 3.3 indicates more consumption smoothing at low income levels in 2009/10 than in 2004/5, which could be expected in the context of changes in the composition of the lowest income decile in 2009/10. This issue is now explored in more detail.

Table 3.4 presents a breakdown by family type of the lowest income decile in each survey year, compared to the population breakdown for each type. There is a large composition change in the lowest income decile between the two years. The single working age adult category comprises 20.7% of low income households in 2009/10 compared to 4% in 2004/5. They have mainly replaced the single adult with children or lone parent category, who have a lower share

of low incomes despite having an increasing population share. Single men were particularly affected by unemployment in the early stages of the "Great Recession" in 2008/9, which contributes to this outcome. Also, income poverty rates fell dramatically for the lone parent group in SILC[7] over the same period due to a fall in median incomes and large welfare increases received in 2005-07, which partly explains the composition change, in addition to large welfare increases received by lone parents in 2005-7. Single adults are now the most overrepresented group in low incomes at 6.1 percentage points above their population share. The shares of lone parents and single adults aged 65+ changed markedly between 2005/5 and 2009/10, reflecting long term population trends. The 2009/10 HBS[8] results also showed an increased presence of transitory low incomes in the bottom decile. 14% of people in the bottom decile of gross household incomes were in third level or post secondary education compared to 6% in 2004/5, while 19% compared to 7% were unemployed. This suggests that more people who had experienced a recent income shock were at the bottom of the income distribution in 2009/10 and the main category they replaced included people who were on predictable low incomes.

Table 3.4 Composition of lowest income decile and population shares, 2004/5 and 2009/10

Household type	Low income decile share in 2004/5 %	Population share 2004/5 %	Low income decile share in 2009/10 %	Population share 2009/10 %
1 adult aged 14-64	4.0	13.0	20.7	14.6
1 adult aged 65+	5.1	13.2	7.6	7.6
Married couple only	17.5	19.9	10.7	15.6
Married couple 1 child	4.2	6.8	6.0	5.6
Married couple 2 children	6.4	9.3	11.1	7.6
Married couple 3+ children	13.1	8.2	9.3	7.1
Single adult with children	21.6	3.2	8.5	7.3
All other households	28.1	26.5	26.1	34.6
Total	100.0	100.0	100.0	100.0

To examine the issue of consumption smoothing by family type in more detail, I now look at the extent to which low income and low expenditure are connected, by family type. The analysis in Table 3.5 considers those households who are located in the bottom decile of the income distribution and reports the proportion of these households that are also located in the

[7] https://www.cso.ie/en/media/csoie/releasespublications/documents/silc/2011/silc_2011.pdf
[8] https://www.cso.ie/en/media/csoie/releasespublications/documents/housing/2010/full.pdf

bottom decile of the expenditure distribution. Table 3.5 shows a stronger overlap in 'low incomes/low expenditure' in 2004/5 than in 2009/10 (58.8% compared to 30.3%) based on the conditional distribution of low expenditure for the lowest decile of population household income. This indicates much higher levels of dissavings at the bottom of the income distribution in 2009/10, which is in line with the theory for a recessionary period. This is also clearly reflected in the comparison of expenditure given incomes in Figure 3.3. Single working age adults have the smallest 'low/low' overlap in both years, which indicates that they can either draw from other resources for current consumption or see their current income situation as being temporary. This suggests a possible transitory income shock, where households are smoothing their consumption levels in anticipation of higher permanent income levels. Lone parents had the highest 'low/low' overlap in 2004/5 (73.6%), which suggests low permanent income. While their overlap is in line with the population average in 2009/10, the composition shift in Table 3.4 makes it difficult to draw conclusions about the change in lone parent dissavings since the group in the lowest income decile is now much smaller. The combined effect of the composition and overlap shifts is that a single adult group that tends towards overspending and consumption smoothing has replaced a lone parent group with low dissavings in the bottom decile of incomes.

Table 3.5 Proportion of people in lowest income decile also in lowest expenditure decile, 2004/5 and 2009/10

Household type	'Low/low' 2004/5 %	'Low/low' 2009/10 %
1 adult aged 14-64	10.5	22.9
1 adult aged 65+	13.0	40.1
Married couple only	56.6	28.5
Married couple 1 child	45.7	33.5
Married couple 2 children	48.7	33.5
Married couple 3+ children	62.2	42.0
Single adult with children	73.6	30.6
All other households	66.5	27.5
Total	58.8	30.3

Slesnick (1993) identified the 'permanent income poor' based on assets, food shares and dissavings. On the basis of dissavings reported in Table 3.5, a group who appeared to be permanent income poor have reduced in prominence in 2009/10, while there are more

transitory income poor. Single men[9] were particularly hard hit by the "Great Recession" as they were overrepresented in construction jobs. The unemployment rate of this cohort increased from 15.0% to 21.4% from Q2 2009 to Q2 2011, compared to an increase from 12.3 to 14.6% overall. For other family types, there is some indication of increased consumption smoothing as all apart from single adults aged 65+ have a lower overlap with low expenditures in 2009/10. However, without an accompanying composition shift into the lowest income decile, this is not particularly convincing evidence of consumption smoothing during an income shock.

The Blundell et al. (2016) diagnosis of family labour supply insurance against income shocks fits well with the limited data presented in Table 3.4 as Irish families in the Married Couple categories appear to have insured themselves against very low incomes. Their share of the bottom income decile shows little change between 2004/5 and 2009/10 and is similar to their population shares in both years There is much broader evidence of insurance in incomes and labour market data for families. Income poverty was relatively flat throughout the "Great Recession" period, with gross household work income – which reflect earnings of all family members - showing remarkable levels of stability after 2008[10]. There were significant shifts in favour of female employment and involuntary male part time employment[11] to maintain household income levels. However, panel data are required to fully explore labour, income and consumption dynamics (frequently mentioned in the literature), so insights in this Chapter are limited to income, expenditure and consumption changes between three five-year cohorts. Income dynamics for households are examined in more detail in Chapter 6 using panel data.

3.5.2 Distribution of income, expenditure and estimated consumption

My analysis to this point is based on expenditure data from the HBS. I now compare consumption to expenditure and income as it presents a more accurate and methodologically sound measure of the resources available to households. The Permanent Income/Life Cycle Hypothesis proposes that where households can consumption smooth then consumption is a better measure of resources than current income. The first Irish set of household survey consumption estimates are presented in this paper using the approach set out in Section 3.4.1. This section compares the distribution of consumption to the distribution of HBS expenditure on which the new estimates are based. The rationale for this approach is to examine whether

[9] http://www.cso.ie/en/releasesandpublications/er/qnhs-fu/qnhshouseholdsandfamilyunitsq22015/
[10] https://www.cso.ie/en/media/csoie/releasespublications/documents/silc/2011/silc_2011.pdf
[11] https://www.cso.ie/en/releasesandpublications/er/qnhs-es/qnhsemploymentseriesq12016/

welfare based on the economic concept of consumption is distributed in line with income and/or expenditure.

Figure 3.4 presents the separate distributions of income, expenditure and consumption in 2009/10, following the O'Neill and Sweetman (2001) presentation based on kernel density estimates. Equivalised expenditure and income closely track each other, with a similar single peaked and skewed distribution to that identified for 1987 and 1994 by O'Neill and Sweetman (2001). Expenditure and income diverge at low levels, which has been widely discussed in recent Irish economic commentary (Collins, 2014) and in earlier analysis in Section 3.5.1. The distribution of consumption is much flatter with lower levels of inequality than incomes or expenditure, which was also found in Meyer and Sullivan (2009) and Brewer and O'Dea (2012) and is consistent with the annuity value of expected lifetime wealth model of smoothing. Based on the S80/S20 Income Quintile Share Ratio, the top quintiles of equivalised income and expenditure had 6.3 and 6.2 times the mass of the bottom quintile in 2009/10 respectively, while the ratio for equivalised consumption was 5.2. There are relatively fewer households in the lower tail of the consumption distribution than in income or expenditure, which is also an important welfare measure.

Figure 3.4 2009/10 Distribution of equivalised income, expenditure and consumption

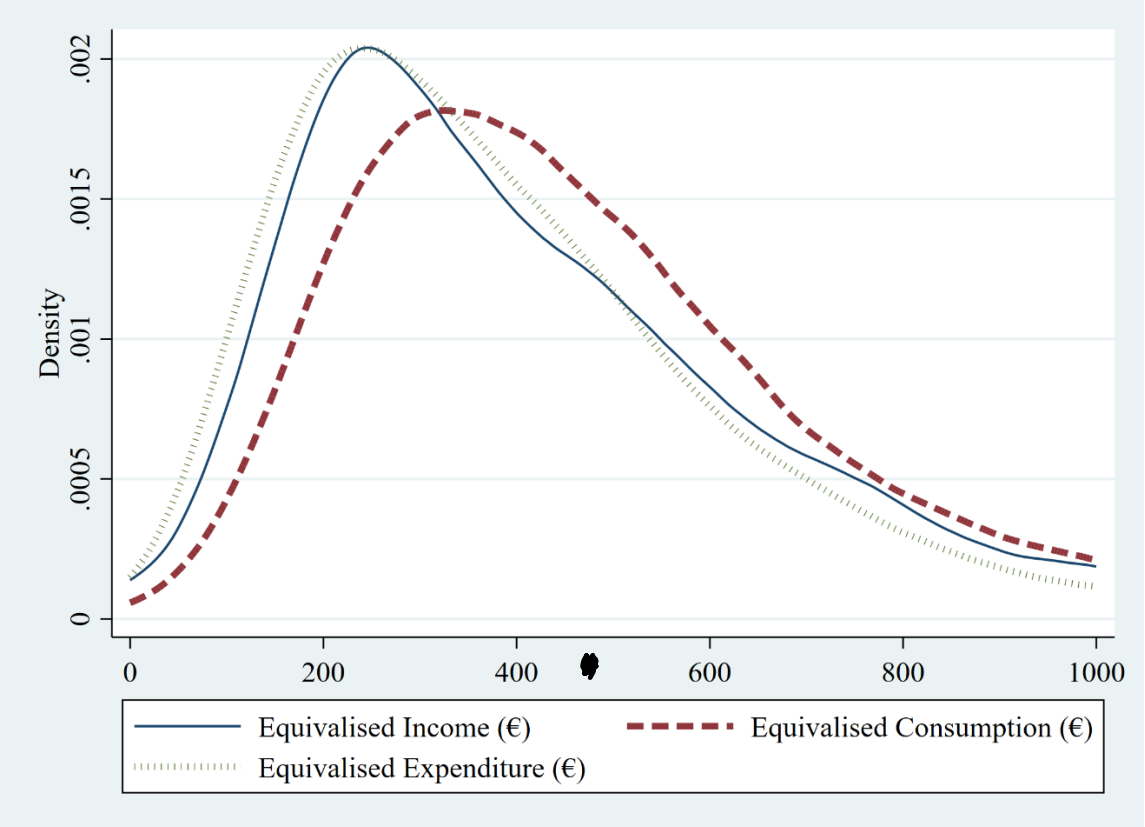

It should be noted that the distribution of income in Figure 3.4 is based on household disposable income, which excludes free medical services and imputed rents. MS/BoD include these consumption items in 'broad income' resource measures. Kernel density presentations of the three resource measures for 1999/2000 and 2004/5 are discussed in Section 3.5.3 and presented in that Section and Appendix 3.

Figure 3.4 shows less mass in the distribution of consumption at low levels compared to expenditure for all households. I now examine the distribution of consumption and expenditure by family type to identify which families are better off on the basis on consumption. Figure 3.5 shows median expenditure and consumption by income decile for all households in 2009/10. There is a consistent 'lift' in living standards from the additional items included in the consumption estimates. This suggests that household living standards are generally higher based on consumption than expenditure.

Figure 3.5 Median 2009/10 expenditure and consumption by income decile (all households)

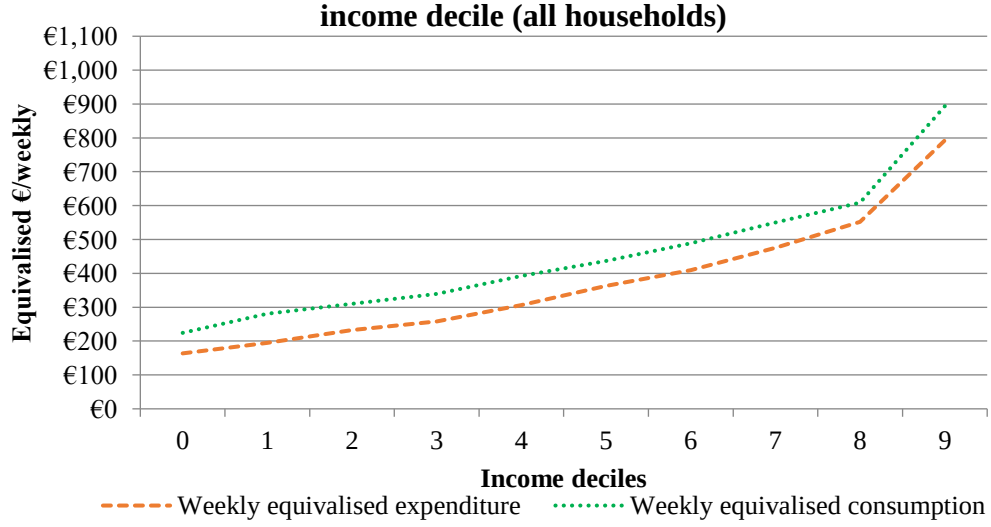

Figure 3.6 shows the same analysis as Figure 3.5 for one adult households aged 65+. Pensioners receive a stronger 'lift' in living standards than the population in general from the methodological adjustments described earlier to calculate consumption. This is mainly due to their tenure status as around 90% are owner occupiers without mortgages and have low expenditure on housing but high consumption. In addition, older people also benefit disproportionately from having almost universal access to medical cards.

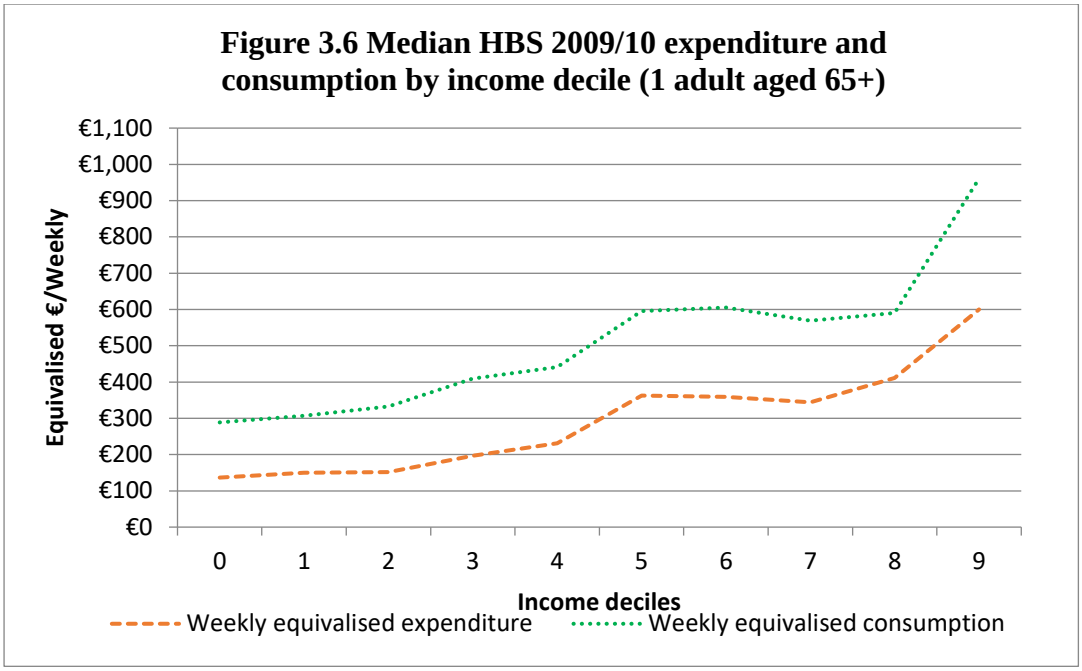

Figure 3.7 shows the same analysis as Figures 3.5 and 3.6 for married couples with 2 children. The 'lift' in living standards from the consumption adjustments is lower than for the overall population, and at the top of the income distribution turns negative.

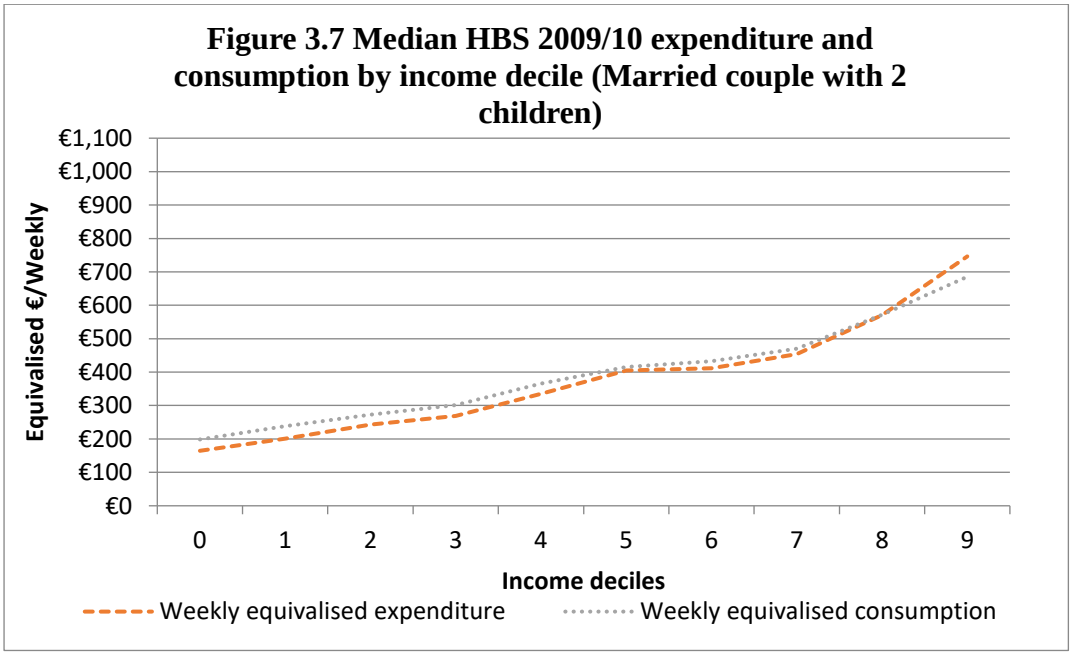

The main reason for this low to negative adjustment is the relative cost of mortgages compared to the service flow. Table 3.1 Product Group 4 'Housing and utilities' shows the increase in HBS housing expenditure from 2005-2010 and the differential with PCE. HBS expenditures are mainly based on mortgage costs for owner occupiers while PCE is based on rental values. This differential is particularly strong for higher income working age groups who tend to have

large mortgages. This effect may have been reversed in recent years due to lower mortgage interest rates for many households and increases in private rents, which are the basis for the imputed rent calculation.

The changing position of lone parent and single adult households stood out in analysis of the lowest income decile in Section 3.5.1. Figures 3.8 and 3.9 show that consumption adjustments appear to be more favourable for low income single adult households in Figure 3.8 than for lone parents in Figure 3.9. This is due to the tenure status of both groups, which is also relevant for Slesnick's (1993) categorisation of the permanent income poor. Lone parents have the lowest level of home ownership among the available household types at 23.4%, while single adult households are broadly in line with the population at 57.5%. Also, roughly half of single adult households are home owners without mortgages, which explains their more favourable consumption outcomes.

Figure 3.8 Median HBS 2009/10 expenditure and consumption by income decile (1 adult aged under 65)

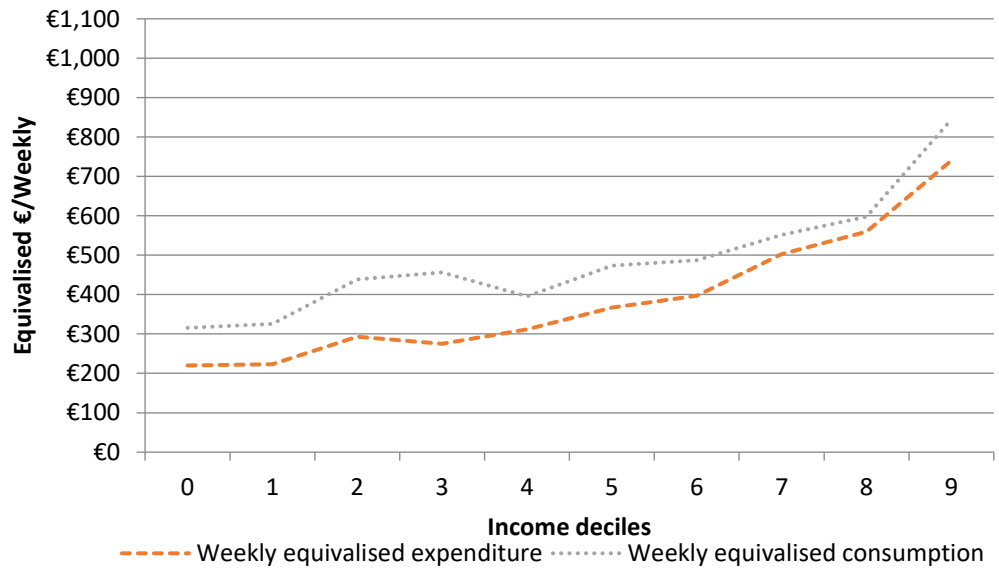

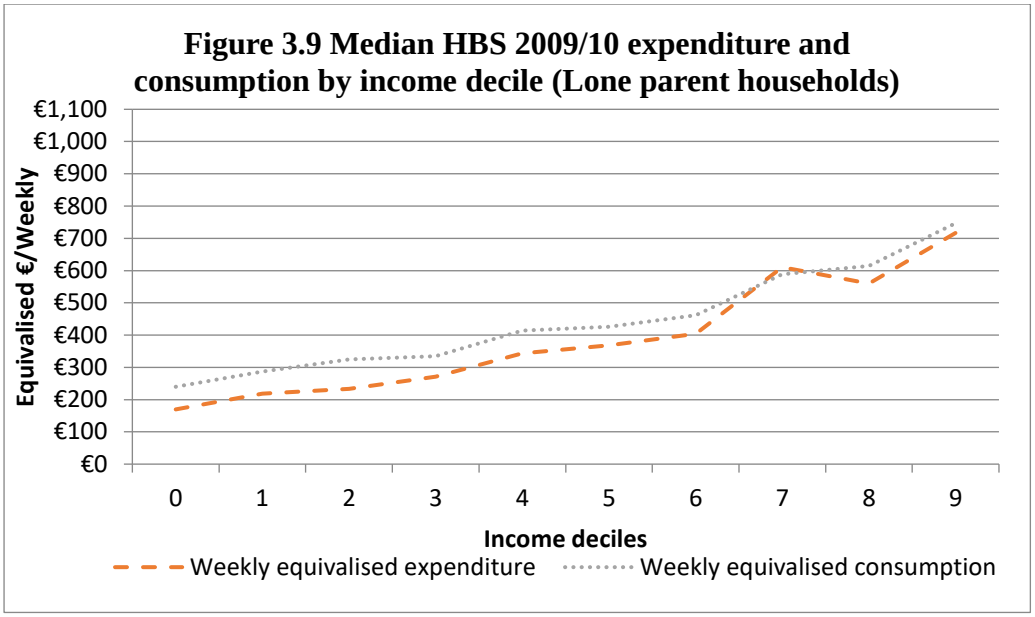

In summary, the distribution of consumption is flatter with lower levels of inequality than incomes or expenditure, which is consistent with UK/US findings and the Permanent Income/Life-Cycle Hypothesis (Figure 3.4). The consumption distribution for pensioners is generally more favourable in than for working age families (Figures 3.5-3.7) which is not consistent with the Life-Cycle hypothesis as discussed earlier. Families with children tend to have smaller differences between expenditure and consumption as they tend to have high mortgage costs relative to market rents.

3.5.3 Income, expenditure and consumption risk-of-poverty analysis

An alternative way to look at differences in the distributions of resources is to compare the groups identified at the bottom of each distribution, following the 'at risk of poverty' approach that is used as a standard welfare measure for incomes. The rationale for this approach is that the group at the bottom of the welfare distribution are of most interest for policy interventions as they are relatively disadvantaged in society. 60% of median equivalised income, expenditure and consumption is used as the threshold below which a household is considered to be 'at risk of poverty' under each resource measure.

Welfare rates were low relative to median incomes for all groups in 1999/2000. Pensioners received large increases before 2004/5 while minimum welfare rates, including lone parents, increased significantly between 2004/5 and 2008. The start of the "Great Recession" is reflected in the 2009/10 data, particularly in a much larger unemployed category. Rents increased between 1999/2000 and 2004/5 and fell subsequently, which affects the consumption

measure particularly. All of these factors feed into the groups identified as poor over time and also in the comparison of outcomes between income, expenditure and consumption.

Table 3.6 shows the size of the group in income, expenditure and consumption 'at risk of poverty' in each survey year in the second last row ('Total'). There is a lower share of all households in consumption poverty each year than in income or expenditure poverty, which is consistent with Figure 3.4 and the theory. Poverty rates by household type vary significantly however.

Table 3.6 Income, expenditure and consumption 'At risk of poverty'* by survey year

Household type	1999/2000 survey			2004/2005 survey			2009/2010 survey		
	Income AROP %	Exp. AROP %	Cons. AROP %	Income AROP %	Exp. AROP %	Cons. AROP %	Income AROP %	Exp. AROP %	Cons. AROP %
1 adult aged 14-64	19.1	15.5	8.0	8.1	9.3	1.9	20.5	18.9	5.0
1 adult aged 65+	27.2	42.2	12.9	5.4	30.0	4.6	26.5	46.3	17.3
Married couple only	12.6	16.0	10.4	17.6	17.4	12.6	24.0	25.5	19.5
Married couple 1 child	16.8	12.6	18.9	10.0	9.1	13.2	11.8	11.3	14.4
Married couple 2 children	19.9	16.6	19.8	12.8	10.6	16.8	15.5	13.2	18.3
Married couple 3+ children	19.6	19.2	26.1	20.5	18.5	25.1	28.2	25.9	33.2
Single adult with children	23.1	23.7	16.1	56.3	50.5	55.3	23.9	23.4	22.6
All other households	10.4	10.2	9.2	14.2	15.3	14.8	10.0	12.8	11.9
Total	16.0	16.6	12.6	14.4	17.4	13.6	18.1	20.8	16.2
of which; Married with children	*18.9*	*16.4*	*21.8*	*14.7*	*12.9*	*18.6*	*18.8*	*17.0*	*22.3*

***Based on 60% of median equivalised income, expenditure and consumption**

Older households (the '1 adult aged 65+' group and a large proportion of the 'Married couple only' category since couples aged over 65 are in this category) have a higher risk of expenditure than consumption poverty for reasons discussed earlier. While income poverty rates for 1 adult aged 65+ are quite volatile, this group is becoming increasingly unrepresentative of older households due to improvements in joint life expectancies. After decades of minimal improvements to male life expectancy, life expectancy at age 60 increased by 3.1 years from 1996-2006 (2.5 years for females). 1 adult aged 65+ are a consistently high-risk group for expenditure poverty over the three periods, which further indicates more frugal lifestyles for older pensioners. Married couples with children (last row) stand out as the high-risk group for consumption in all years due to relatively high mortgage costs, also as discussed earlier. This

group were 6.1 percentage points above average for consumption poverty compared to 5 p.p. in 2004/5. Consumption poverty rates for married couples with children in 1999/2000 are much higher than in 2009/10 (9.2 compared to 6.1 p.p. above average) which suggests a much more unequal distribution of housing costs for families that is consistent with younger family formation. Household incomes are lower at younger ages while mortgage costs tend to be higher as a proportion of incomes at the outset of home ownership.

The standout feature in the 2004/5 data is the higher percentage of the population in expenditure poverty than in the other two measures. This issue is examined next in conjunction with kernel density estimates. The main poverty risk group on all measures in 2004/5 were lone parents, which was an issue of public policy concern at the time.

The relatively high percentage of the population in expenditure poverty compared to income poverty in 2004/5 identified in Table 3.6 is either due to higher expenditure patterns at median levels or downward shifts in the bottom of the expenditure distribution. To illustrate the impact of median growth, increases in median incomes drove much of the increase in the risk of poverty in the late 1990s in Ireland as social welfare incomes increased by about half of median income levels. Table 3.7 shows that median increases similarly drove higher expenditure poverty rates, since overspending (expenditure divided by income) is higher in Deciles 4-6 in particular. We also see that overspending is less prominent in the lowest household income decile in 2004/5 than in 2009/10, for reasons discussed in Section 5.1. 'Wealth effects' due to high property prices during the boom may have contributed to high 2004/5 spending levels. Housing wealth in particular is indirectly measured through consumption as it reflects the economic value of housing to owner occupiers.

Table 3.7 Ratio of expenditure to income for gross household income deciles by survey year

	Decile 1 Ratio	Decile 2 Ratio	Decile 3 Ratio	Decile 4 Ratio	Decile 5 Ratio	Decile 6 Ratio	All Ratio
Expenditure/Income 2004/5	1.37	1.25	1.15	1.15	1.06	1.06	0.93
Expenditure/Income 2009/10	1.89	1.28	1.11	1.05	0.98	0.97	0.92

Source;https://www.cso.ie/en/media/csoie/releasespublications/documents/housing/2010/full.pdf

Figure 3.10 presents kernel density estimates for the three resource measures in 2004/5, which shows the full distribution of the three resource measures. Consumption is much flatter in 2004/5 than in Figure 3.4 for 2009/10 as imputed rents contributed more strongly to household resources in 2004/5 than in 2009/10 since mortgages were lower and rents were higher. Figure 3.12 in Appendix 3 gives the same presentation for 1999/2000, and it shows a much clearer peak in all three resource measures, similar to 2009/10.

Figure 3.10 2004/5 Distribution of equivalised income, expenditure and consumption

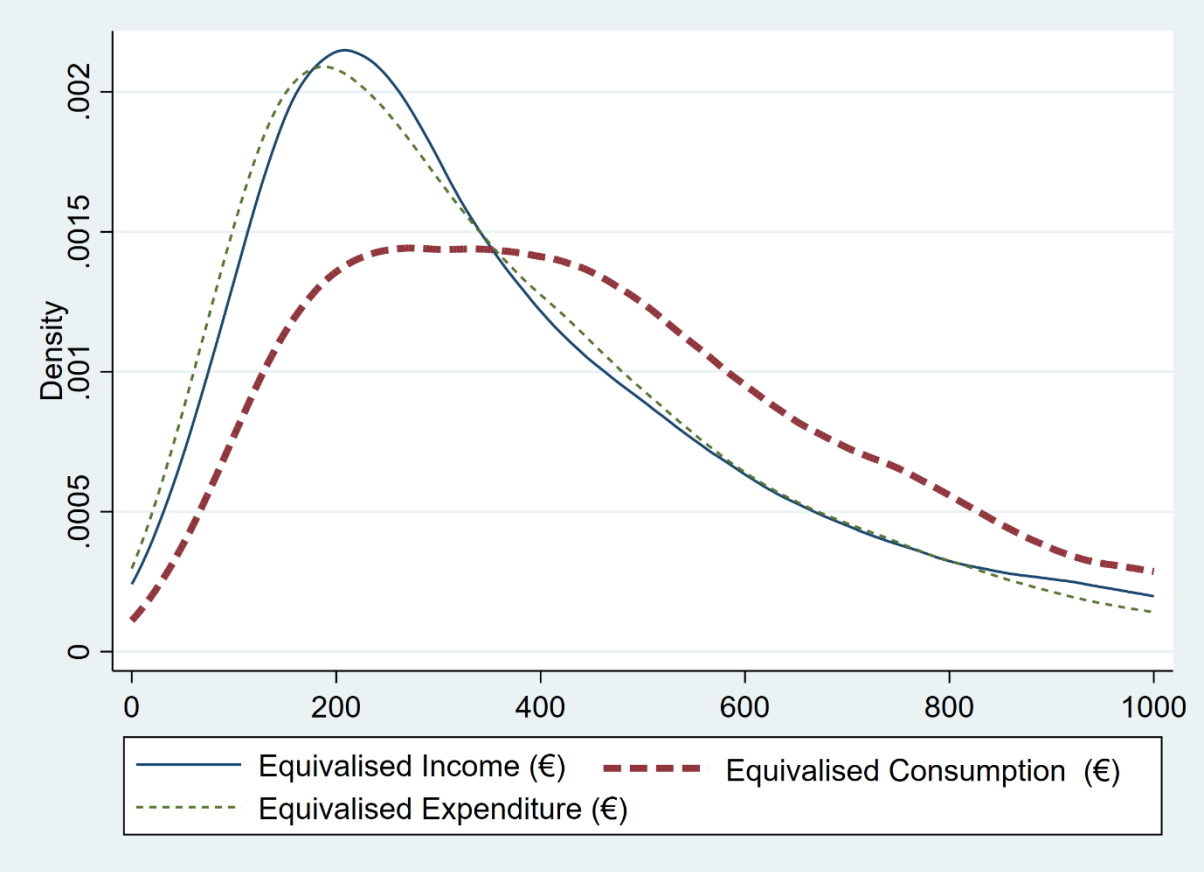

There are different findings in the at-risk-of-poverty analysis in this paper to the Irish literature for poverty identification, which was based on incomes only, when different resource measures are used in addition to incomes. Callan, Bercholz and Walsh (2018) summarised most of the

previous research and found that, despite recessions, household incomes have grown strongly in Ireland in the last 30 years with stability across the income distribution. Consumption identifies working age families with children as an at-risk group not identified by other resource measures in all years, due to high housing costs. Expenditure in 2004/5 was out of step with income and consumption based on the poverty and kernel density estimates presented in Table 3.6 and Figure 3.10. This may be related to wealth effects from housing which are indirectly picked up in the consumption measure. However, as with the conclusions on Permanent Income in Section 3.5.1, panel data, ideally including wealth/financial balances as well as income, expenditure and consumption, is needed to explore the dynamics of how overspending was financed at micro level.

3.6 Conclusions and Discussion

This paper uses empirical strategies from US and UK research that are well established to assess the extent of measurement error in the Irish Household Budget Survey (HBS). Measurement error is found to be lower than in similar US and UK surveys for income and expenditure. The survey process for the HBS changed to collection from administrative sources for welfare incomes in 2009/10 from direct responses, which facilitates a 'before and after' comparison that adds to the current literature on mismeasurement of low incomes due to respondent error compared to unrealistic consumption levels. However, there is no apparent mismeasurement of low incomes in HBS for 2004/5, which suggests that the relatively simple Irish benefits system is a possible reason for better measurement of low incomes since respondents accurately reported their incomes. This approach also identifies more consumption smoothing at the very lowest income levels in 2009/10 than in 2004/5, a finding that supports the Permanent Income Hypothesis. A group who were 'transitory income poor' (mainly unemployed single adults, who showed evidence of consumption smoothing) replaced a group who were 'permanent income poor' (lone parents) at the bottom of the income distribution, which is in line with expected outcomes from the theory in a recession period.

Irish household consumption estimates based on economic concepts from the National Accounts are constructed in this paper for the first time using HBS expenditure data. Comparisons of the distribution of income, expenditure and consumption reveal some new insights into life cycle and household resources trends in Ireland for the period from 1999 to 2010. Consumption is flatter than the other two measures in all years, which is consistent with smoothing. Consumption identifies different risk groups based on 'at risk of poverty' in all

years to the other resource measures, mainly as mortgage costs are replaced by rental values in consumption. Working age families with children are a risk group in all years due to high mortgage costs, while pensioners are lower risk than for incomes or expenditure since they rarely have mortgages. For expenditure, risk groups are similar to incomes in all years, but the at risk of poverty rates were significantly higher for expenditure than incomes in 2004/5. This was driven by higher expenditure for median groups. If people at the bottom of the distribution assessed their living standards using expenditure instead of incomes they may have felt relatively more disadvantaged. On the basis of evidence of consumption smoothing and differences in the risk groups identified, in addition to the theoretical justification from Chapter 2, it is possible that consumption was a better measure of household resources than income or expenditure during the period of analysis.

Appendix 3

Figure 3.11 Median expenditure by income, and median income by expenditure, HBS 2009-10 including Cumulative Density Functions

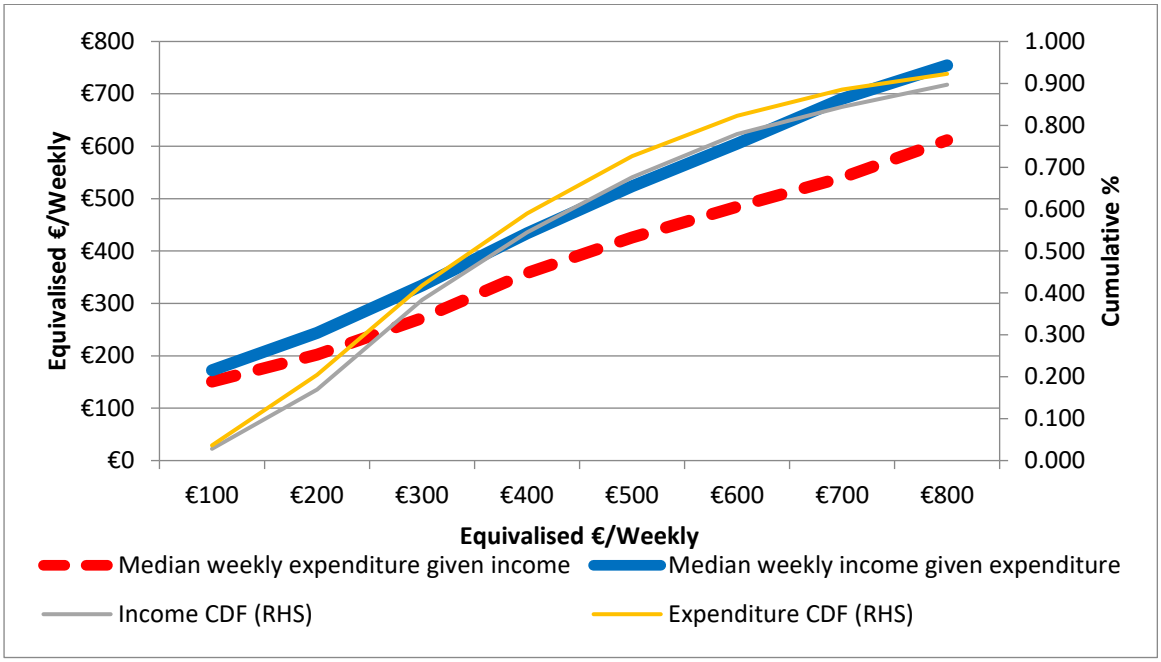

Figure 3.12 1999/2000 Distribution of equivalised income, expenditure and consumption

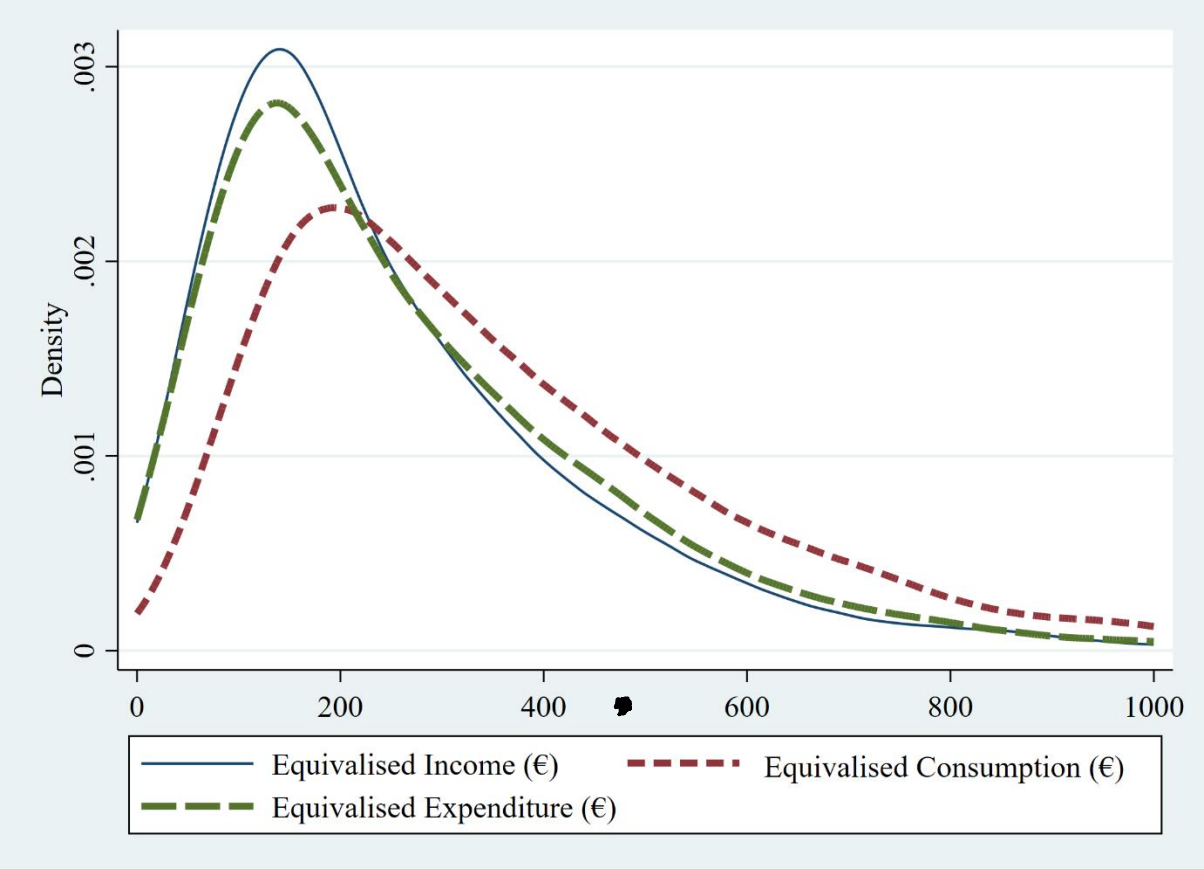

4. Retirement Consumption Puzzle

4.1 Introduction

Having examined the data and noted the importance of consumption in understanding inequality and poverty, I now turn to assessing an apparent challenge for the PIH. Empirical research in the 1980s was mainly based on the 'certainty equivalent' model. Under this model, consumers change their consumption plans in response to permanent, unanticipated income shocks. Consumers also behave as though future income is predictable at its expected value, and they thus have a benign attitude to risk. There is substantial literature on the 'retirement consumption puzzle', which is an unexpected dip in consumption at retirement age common to many countries. This is inconsistent with the 'certainty equivalent' model, since retirement age should have no impact on consumption planning. I will examine this 'puzzle' in the Irish context – whether the retirement dip occurs, and, if so, the reasons for it.

The current Irish literature on the 'retirement consumption puzzle' (Redmond and McGuinness, 2020) finds two product groups with significant reductions in expenditure at retirement in a selection of seven Household Budget Survey (HBS) product groups representing a third of total expenditure. Food expenditure falls and is concentrated in meals out, which is in line with the life cycle as it is partly a work-related expenditure. Services also falls which is more concerning as it could represent an enforced reduction in expenditure due to lower than expected retirement incomes. I expand the analysis in Redmond and McGuinness (2022) in a number of ways. Firstly, rather than focusing on a subset of products, I examine all HBS products. Secondly, I look at consumption rather than expenditure, using the approach described in Chapter 3 to convert expenditure to consumption flow values. Analysis of consumption in the literature finds that it is less sensitive to the retirement transition than expenditure, particularly due to stability of the flow value from owner occupied housing (Fisher et al., 2008). Nevertheless, I still find a reduction in consumption at retirement age on this basis. Looking across products, I find a similar result on food to Redmond and McGuinness, and furthermore that most of the overall reduction in consumption in retirement is in the services category. Also, my approach allows for international comparisons as it covers all household consumption. The statistically significant fall in services found for Ireland is not found in the comparable international studies and could be considered as a failure of the 'certainty equivalent' model. While pensioners maintained their relative living standards based on

incomes during the period from 1994 to 2009, they may be falling behind in these types of discretionary consumption.

4.2 Literature review

As discussed in Chapter 2, the Permanent Income/Life-Cycle Hypothesis proposes that permanent income, which is the annuity value of lifetime resources, or consumption (which depends only on permanent income) is a more appropriate measure of household resources than current income. Under the theory, households smooth their lifetime consumption against short term income shocks or 'transitory income' shifts. Insurance against shocks is usually facilitated through savings. Over the longer term, households maintain 'life-cycle' consumption through a similar savings process.

There is widespread empirical evidence against the Permanent Income/Life-Cycle Hypothesis, based on patterns observed in data where consumption tracks income in the short and long term. Thurow (1969) was concerned that the correlation between income and consumption over a working life did not fit with the life-cycle theory. In cross section US data, both income and consumption had a similar inverted U shape over the life cycle, which has also been widely observed in many other countries and datasets. Thurow (1969) believed that this was because households were liquidity constrained, and thus could not fund their desired consumption path. Nagatani (1972) argued that households were prudent and save more when faced with riskier future income as described in Chapter 2. These models add a precautionary motive for saving in addition to the consumption smoothing motive. Increased uncertainty related to future incomes induces precautionary savings by prudent households. Consumption is therefore lower in these models than under certainty equivalence, and also more closely tracks income (Caballero, 1990). Heckman (1974) suggested that the inverted U shape was due to labour supply, which affects consumption through earnings and income. He gave examples where labour and consumption were complements which further reinforces the relationship. The most long standing and obvious explanation is that the presence of children determines the U-shape (Tobin, 1967). Economists have responded to these findings by proposing variations on the life cycle model which better reflect conditions facing households in the 'real world'.

Tests of the certainty equivalent model started with Hall (1978), who proposed that consumption reflects expected future income (i.e., certainty equivalence) and thus only responds to unexpected events and is unpredictable. He found empirical evidence in favour of his hypothesis in the US National Accounts with the change in consumption being unaffected

by lagged incomes. Flavin (1981) found contradictory evidence from US macroeconomic data. In particular, consumption was 'excessively sensitive' to predictable changes in incomes, which is a failure of the Permanent Income Hypothesis (PIH). These early US tests of the PIH were particularly sensitive to timing and the number of lags considered. Hall and Mishkin (1982) used food expenditure from the US Panel Survey of Income Dynamics (PSID) and found that the proportion of consumers where consumption closely tracks income was 20 per cent. This group was identified as 'rule of thumb' consumers who consumed a fixed proportion of income. For the rest of the panel, consumption was in line with the PIH. They found that food consumption was positively related to predictable changes in income, but with a planning horizon of only three to four years into the future. These results confirm Milton Friedman's (1957) initial view that "the permanent income component is not to be regarded as expected lifetime earnings. . . . It is to be interpreted as the mean income at any age regarded as permanent by the consumer unit in question, which in turn depends on its horizon and foresightedness". In addition to the excess sensitivity of consumption to income, consumption also appears to be too smooth, which suggests that consumers are not taking new information into account in consumption decisions (Flavin, 1993).

Hayashi (1985) directly identified liquidity constraints in a single year of US cross section data. He split the sample into a group that are possibly liquidity constrained and a group with high savings rates who are not. Consumption is regressed on assets, income, age, and interaction terms, using a Tobit model for the unconstrained group. A Tobit model was used as it designed to estimate linear relationships between variables where there is censoring in the dependent variable. This gives an estimate of 'desired consumption' which can be compared to the Ordinary Least Squares average estimate for the entire sample. He found that the Tobit estimate for high savings households overpredicted consumption for the whole sample, and thus inferred that low savings households were unable to consume as much as desired. Campbell and Mankiw (1989) found that the elasticity of intertemporal substitution was very low (i.e. low sensitivity of the growth rate of consumption to the real interest rate) which is another failure of the PIH. Based on this evidence, they found that a mix of 50 per cent 'rule of thumb' consumers and 50 per cent forward planners best fits US macroeconomic data. Carroll and Summers (1991) presented evidence against the 'rule of thumb' based on savings data, and favoured liquidity constraints and prudence as the likely explanations for excessive sensitivity to income.

The effect of family composition on consumption has also been widely examined. Attanasio and Weber (1995) found consumption growth to be consistent with the life-cycle model when family composition and labour supply are taken into account. They analysed Consumer Expenditure Survey data in detail and rejected findings by Hall (1988) and Campbell and Mankiw (1989) against the life-cycle model based on aggregation bias in macroeconomic data. For example, aggregation hides savings by age, which is necessary information to test the PIH. They also found food to be a poor proxy for total consumption. Attanasio et al. (1999) returned to this topic and found that a precautionary motive is needed in addition to accounting for family size in order to explain the correlation between income and consumption over a life cycle. Villaverde and Krueger (2007) found that 50 per cent of the 'hump' in expenditure in middle age is due to household demographics, though this result depends on the choice of equivalence scale used. The rest of the 'hump' is attributed to factors mentioned earlier – complementarity of labour and consumption and liquidity constraints – as well as 'uninsurable uncertainty' for events such as long-term invalidity or redundancy.

The literature on the 'retirement consumption puzzle' is based on evidence of a substantial fall in consumption in retirement, which could contradict the PIH. Banks et al. (1998) found that British households reduce consumption at retirement age by an amount that is not consistent with the PIH/life-cycle model. They also found that either people underestimate their future retirement income or retire early due to ill health. Early retirements are an unanticipated shock and reduced consumption would not constitute a failure of the PIH on this basis. Using US panel data, Bernheim et al. (2001) found that poor retirement planning based on a 'rule of thumb' approach was the main reason for the fall in consumption. Using UK panel data, Smith (2006) found unanticipated early retirement to be the main reason, which is consistent with the PIH. An alternative specification for the puzzle was developed by Battisin et al. (2009). They used a regression discontinuity approach based on observed consumption either side of the retirement event. They found a smaller drop in consumption in retirement of 9.8 per cent compared to 15 per cent in Banks et al./Bernheim et al. and that the drop was due to lower work-related expenses.

Much of the subsequent research to Banks et al. (1998) and Bernheim et al. (2001) has sought to explain the puzzle in the context of either better coverage of consumption – both studies only examined food expenditure - or distributional results. Panel data has been widely used to examine distributional results (e.g., Hurd and Rohwedder, 2008) and generally finds the 'puzzle' to be concentrated in low income or wealth groups. Aguila et al. (2011) found evidence

in the US that food expenditure falls unexpectedly at retirement but overall non-durable expenditure does not change significantly. Extensions to the Banks et al. (1998) approach include analysis of surveys/panels with recall questions for the fall in consumption or actual retirement age, and also surveys with planned retirement age questions (Miniaci et al., 2003, Hurd and Rohwedder, 2005 and Haider and Stephens, 2007). Further variations include expanded definitions to capture consumption service flows (e.g. Fisher et al., 2008), analysis of scanner data for evidence of 'shopping around' by older people (Aguiar and Hurst, 2005) and life cycle model simulations (e.g., Blau, 2008). Most find a smaller fall in consumption at retirement than Banks et al. (1998) and Bernheim et al. (2001) and explain this mainly within the PIH/life-cycle context since reductions in consumption at retirement are in work-related expenditures or are concentrated in low income groups who may have suffered income shocks. Hurst (2008) summarised critiques of the Banks et al. (1998) and Bernheim et al. (2001) findings as follows. The declines in expenditure are limited to work-related items - mainly transport, clothing and food. The decline in food expenditure does not reflect lower food consumption, as retirees have more time to 'shop around' and prepare meals. Analysis based on more comprehensive measures of consumption (e.g., Fisher et al., 2008; Aguiar and Hurst 2008), including housing service flows in particular, shows no significant decline at retirement. Finally, the fall in expenditure is concentrated in the bottom income quartile who are likely to have experienced involuntary retirement. These critiques of the early findings on the puzzle inform the rest of my analysis in this chapter. I firstly examine if conversion to consumption values diminishes the change in expenditure at retirement. I subsequently look at product level detail to assess whether falls are consistent with the PIH.

4.3 Data

The anonymised HBS datasets lodged with the Irish Social Sciences Data Archive (ISSDA) for the survey years 1994/95, 1999/2000, 2004/5 and 2009/2010 are used for this analysis supplemented with details on the age of the Household Reference Person in each dataset on request from CSO. The main purpose of the HBS is measurement of household expenditure for price index weights, which requires a very detailed categorisation of expenditure by households. The HBS also includes detailed information on income receipts by source and household composition for nationally representative samples of Irish households. For incomes, household disposable incomes are used in the analysis in this paper which take account of taxes and social welfare incomes. Consumption is constructed in this paper based on HBS expenditure and additional information from the National Accounts and Census of Population.

The adjustments for consumption are described in detail in Chapter 3. Expenditure in HBS is based on a two-week diary of household purchases with a lookback for large occasional purchases over the previous year. Equivalisation[12] accounts for much of the anticipated change in consumption in the life-cycle theory due to changes in family composition (Attanasio and Weber, 1995; Villaverde and Kreuger, 2007). The HBS sample sizes for 1994/5, 1999/2000, 2004/5 and 2009/10 were 7,877, 7,644, 6,870 and 5,891 households respectively, and I will refer to each of these survey years as 1994, 1999, 2004 and 2009 from now on.

While there is a mandatory age of 66 for social insurance-based pensions in Ireland, many people in Ireland retire before that age. Some are ill-health retirees, who are mainly paid a social insurance Invalidity Pension, but most are voluntary and are either funded by their employer or by employee savings, such as Additional Voluntary Contributions. Some public service jobs also have earlier retirement ages, such as teachers, police and the army. Based on this range of outcomes, there is no retirement 'cliff' in aggregate incomes at age 66 in Ireland, and a more gradual fall over a range of ages should be expected. As I use increasing age as a proxy for retirement in the following analysis, there is clearly a threat to identification from early retirements since any fall in consumption related to retirement cannot be contained in a chosen window of ages. The difference in average consumption between two samples with Heads of Household aged 56-61 and 62-67 is not statistically significant at the 95 per cent confidence level, which may be affected by retirements between ages 56-61. Also, retirement ages are likely to change over time which affects comparisons between cohorts.

4.4 Consumption by age for four HBS cross sections (1994, 1999, 2004, 2009)

The analysis presented in this paper is based on four cross sections of HBS data to find evidence of a dip in consumption at retirement age. I will firstly present results for each cross section of HBS respondents to give an indication of how consumption changes as the respondent or Head of Household (HoH) age increases. Figure 4.1 shows that weekly equivalised consumption is higher than equivalised income at all ages after 35 for the 2009 cross section but particularly so after age 50. While expenditure cannot remain higher than income for long periods, consumption also includes a 'flow value' from purchases in previous years where the consumer still benefits from the purchases. The two significant adjustments I make to expenditure for this

[12] Household incomes and consumption are equivalised in this paper based on the National equivalence scale used in ESRI/CSO analysis of SILC incomes, which assigns a weight of 1 for the first adult, 0.66 for the second adult and 0.33 for children aged 14 and under.

purpose are flow values for houses and cars purchased. The flow value for housing is based on current market rents, which increased significantly during most of the 1994-2009 period. Income and consumption are volatile from year to year due to small sample sizes but a drop in consumption levels around retirement age at ages 62-67 is still apparent. There is a substantial increase in household consumption levels from age 45-55. This is also seen in most of the international research and is usually explained by failures of the PIH since equivalisation accounts for changes to household demography (Attanasio et al., 1999; Villaverde and Kreuger, 2007). Figures 4.11-4.13 in Appendix 4 present the same analysis as Figure 4.1 for the earlier cross sections. The shift in levels at ages 45-55 is not as prominent in these years, but otherwise both income and consumption follow the inverted U shape described in the literature in all years.

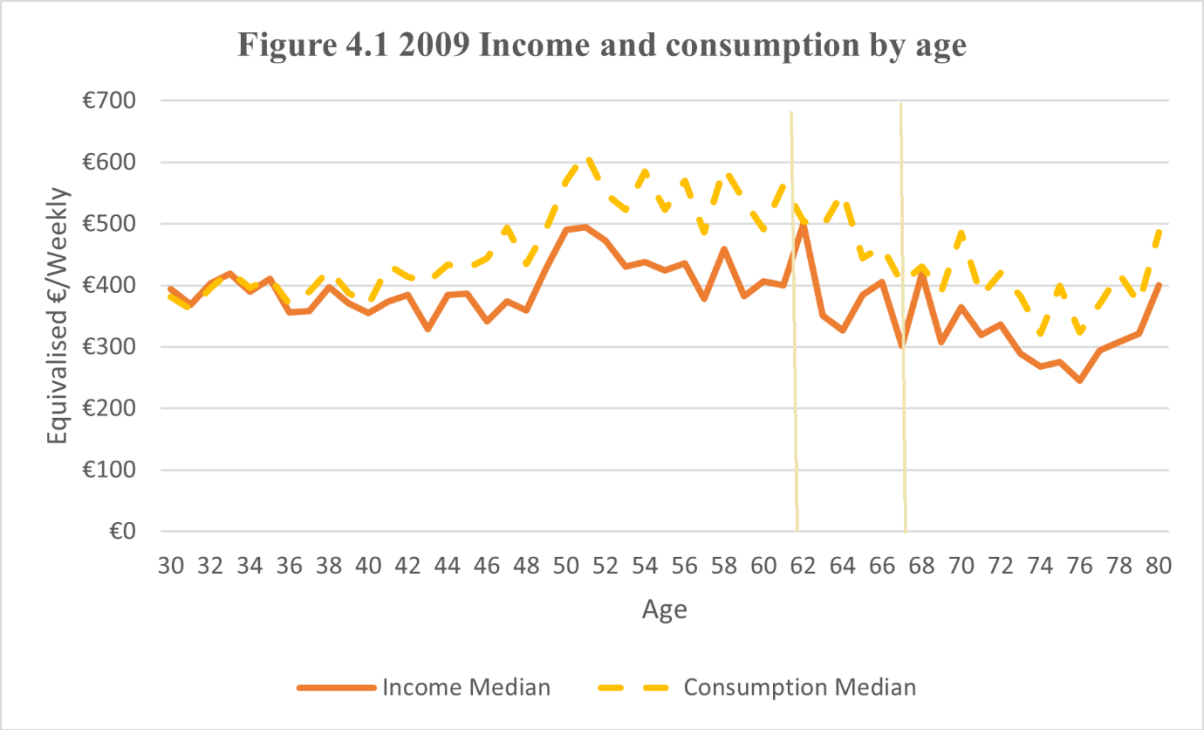

One of the main reasons for the divergence between income and consumption in Figure 4.1 is that median housing consumption[13] continues to increase during these ages despite declining housing expenditure on mortgages. The housing consumption value, or 'imputed rent', reflects economic concepts of consumption which include flow values from durables. Home ownership is part of retirement planning for most Irish people, which is reflected in higher resources of

[13] As described in Chapter 3, the rental value of housing is included in the consumption estimates rather than mortgage outlays. Since older people are more likely to live in bigger houses in well established areas, the age profile of consumption is rising.

older people in the age profile of consumption compared to incomes. This analysis is comparable with results from Fisher et al. (2008) and Aguiar and Hurst (2008). Figure 4.2 below shows that owner occupiers pay down their mortgages at these ages while retaining the service flow or imputed rent from living in their house. The addition to expenditure from the consumption flow value ('Housing consumption adjustment' in Figure 4.2) is lower in ages 30-50, since the service flow replaces mainly mortgage expenditures at these ages.

Figure 4.2 2009 Median equivalised housing consumption and the consumption adjustment by age

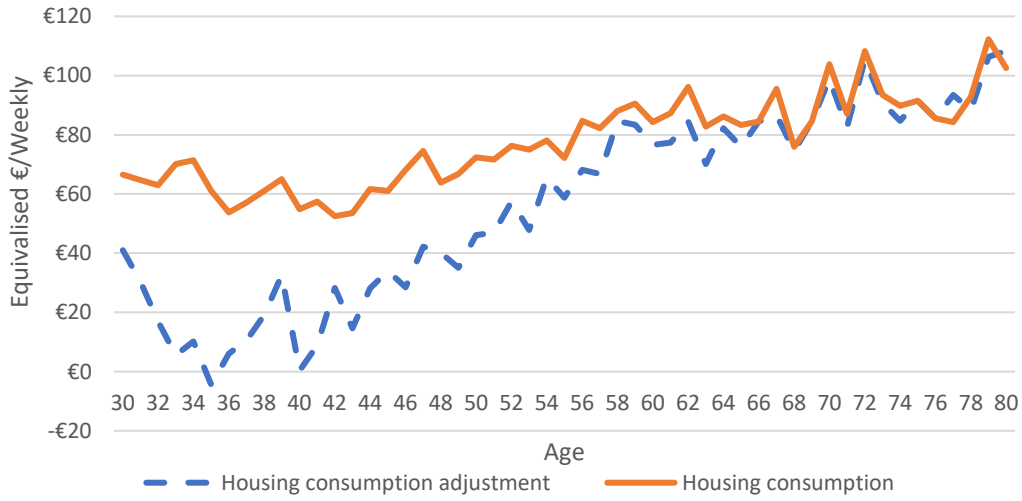

Next I show how income and consumption vary between five year cohorts from working to retirement age to find evidence of consumption smoothing through this transition. Figure 4.3 describes the change in income and consumption levels as age of HoH increases based on five year moving averages of median income and consumption for 2009 households. For example, the change for HoH aged 52 is based on the sum of the five yearly changes for ages 50-54 compared to the sum of the changes over ages 45-49. The rapid increase in levels in Figure 4.3 from ages 45 to 50 is more obvious in the five-year moving average results presented here than in Figure 4.1. Income and consumption increase again after age 77 which is due to higher life expectancy for better off pensioners who tend to survive longer. Consumption is 18 per cent lower for HoH aged 67 than at age 62. By comparison, Banks et al. (1998) found a fall in expenditure at these ages of 15 per cent and Bernheim et al. (2001) found a fall of 14 per cent. Based on a like-for like comparison using consumption instead of expenditure, Aguiar and Hurst (2008) and Fisher et al. (2008) found falls of 5 and 1.2 per cent in total consumption respectively.

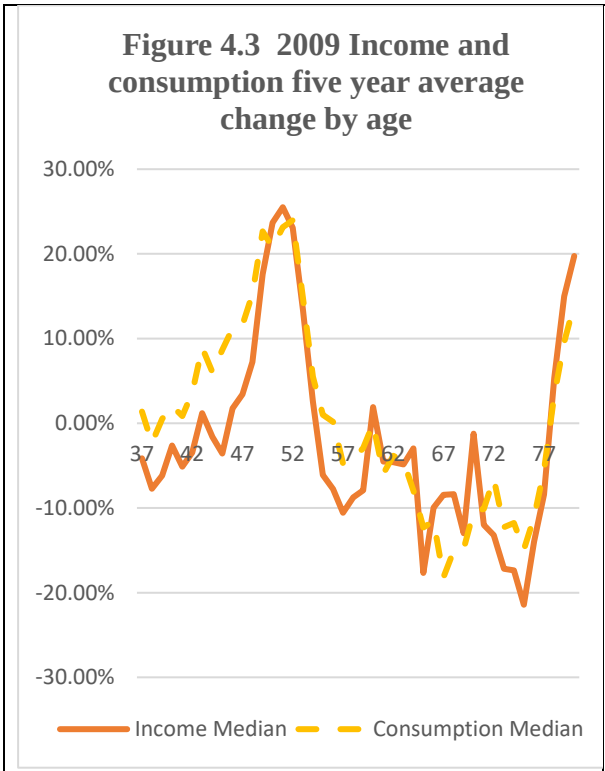

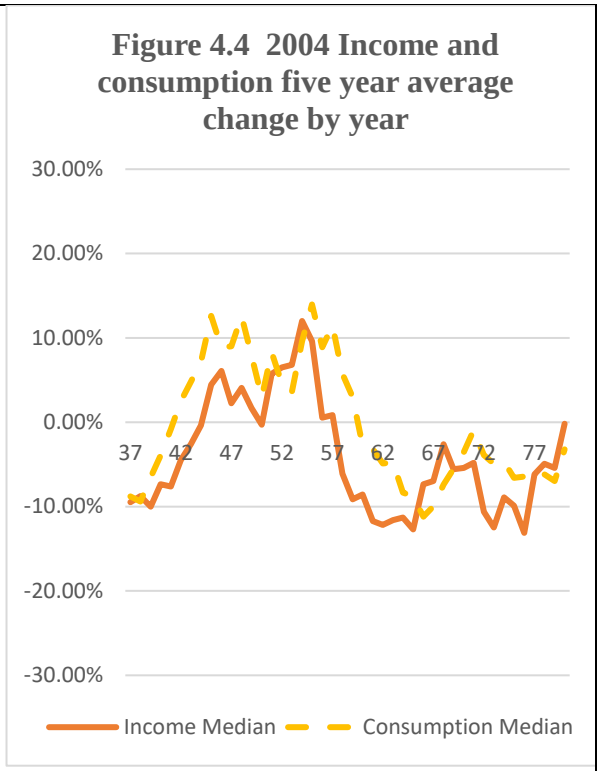

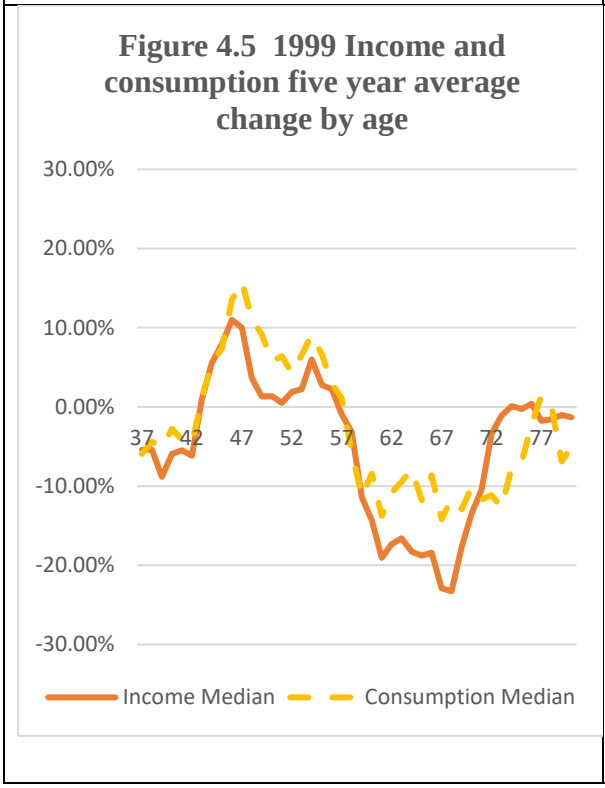

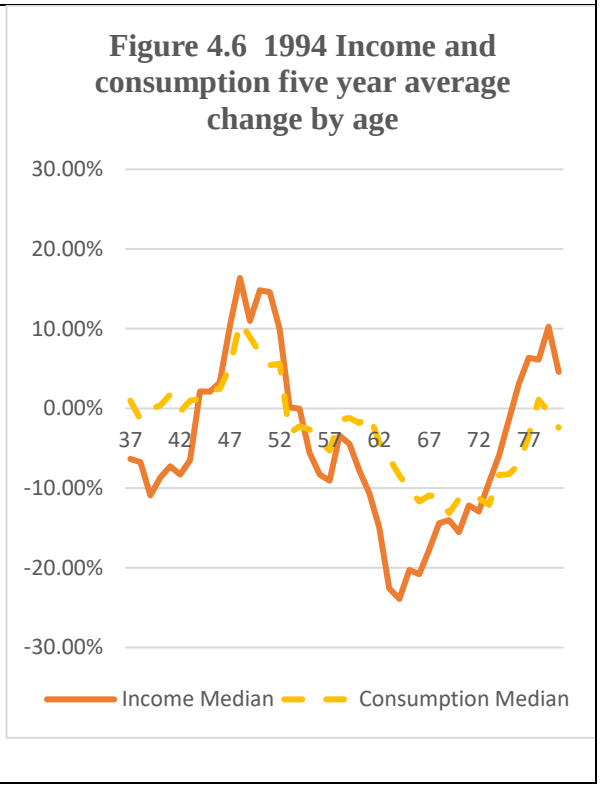

Figure 4.4 presents a similar pattern of falling consumption at ages 62-67 for 2004 respondents, but the reduction is lower at 11 per cent. Imputed rents were stronger in 2004 than in 2009 because of higher market rent levels, which cushioned the overall fall in consumption at these

ages. The fall in incomes by age persists over a much longer period than in 2009, from ages 57-67, which suggests more variation in retirement ages in 2009.

Figure 4.5 presents the pattern of income by age of HoH in the 1999 cross section. The fall in income by age is almost twice as severe as in 2004. This could be because social welfare pensions had fallen significantly behind overall incomes. There is clear evidence of consumption smoothing in 1999 as the fall in consumption with age is around half that of incomes. Figure 4.6 presents the pattern of income and consumption in the 1994 cross section and is almost identical to 1999, with a large fall in retirement incomes by age and evidence of consumption smoothing.

Consumption levels are almost always above income in the data due to the flow value economic concept, where consumption can be derived over many years from expenditure on durable goods. There are very low levels of reported savings in the HBS (income less expenditure), and HBS expenditure is the starting point for the consumption measure in this paper. While it could be expected that households were under pressure in 2009, this picture of low savings holds across all years. The adjustments to derive consumption estimates from HBS expenditure are all positive in aggregate which turns savings negative, apart from the adjustment for expenditure on cars which redistributes expenditure levels among car owners. However, as described in Chapter 3, measurement of consumption and income in HBS stands up well to tests from the literature.

4.5 Consumption at commodity level by age within and between cross sections

The initial evidence presented in Section 4.4 suggests that there is a reduction in household consumption at retirement ages which exceeds that expected under the PIH or the experience in the comparable research. Much of the discussion around the retirement consumption puzzle is concerned with commodity level detail. In particular, both work-related expenses and housing service flows are of interest. Reduced work-related expenses do not imply a failure of the PIH. Also, Fisher et al. (2008) and Aguiar and Hurst (2008) emphasised the importance of housing service flows, since most pensioners achieve part of their consumption smoothing objectives through home ownership. In this section, I look at consumption by commodity groups. The three commodity categories used here are exhaustive and include service flows for housing and vehicles. To give an indication of the relative size of the categories in 2015, food, services and other expenditure comprised 15, 34 and 51 per cent of total expenditure. The most significant sub-categories within each group are highlighted in Table 4.1 below.

Table 4.1 Largest sub-categories within each of the HBS consumption categories, 2015		
Food	***Services****	***Other consumption***
Meals away from home	Insurance	Housing
Meat	Holidays	Transport
Vegetables	Telecommunications	Fuel and light
Fruit and nuts	Medical expenditure	Clothing and footwear
*Private pension contributions are included in this HBS expenditure category but are excluded here		

Food and services[14] consumption are identified separately since they are categories of interest in the research, and the rest of consumption is presented as a residual. The commodity groups are deflated by the most appropriate CPI sub-index in each case.

Figure 4.7 presents 2009 consumption growth by age using the same five-year average approach as Figures 4.3-4.6 for the three broad consumption categories. Consumption on both food and services is noticeably lower for households with a HoH aged 67 than aged 62, which largely explains the fall in consumption of 18 per cent at this age seen in Figure 4.3. Housing is the largest individual component in 'Other Consumption'. The age profile of housing consumption is quite flat as seen in Figure 4.2. This category otherwise shows a slow decline after age 60 as work-related expenditures such as clothing and transport fall gradually over time. The fall in food consumption at retirement is in line with the international literature. For Ireland, Redmond and McGuinness (2020) found this reduction in HBS food expenditure to be concentrated in 'meals away from home' which fits with the PIH, and particularly with life-cycle planning. The services fall was also highlighted by Redmond and McGuinness (2020). However, they only examined specific categories of expenditure amounting to 35 per cent of the HBS total, making their analysis difficult to compare with international results. There is a significantly bigger fall in consumption at retirement found in my analysis than in comparable international research that accounts for service flows, such as Fisher at al. (2008). Also, the fall is concentrated in services which has only been identified as a category of interest in the Irish context. Redmond and McGuinness (2020) particularly focussed on the fall in leisure related activities in this category. They highlighted the social consequences of this in terms of isolation and possibly higher mortality. There are also some sub-categories within the broader category of services in my analysis which are age related, such as life insurance and third level education expenses for children. While results are similar across categories to my analysis, Redmond and McGuinness (2022) used a 'difference in differences' approach to identify the fall in expenditure that is related to retirement by comparing the change in expenditure for a group of

[14] Services expenditure from HBS is adjusted for free medical services provided to medical card holders

people who had retired between ages 55-64 and 65-74 with a group who continued working. This addresses the identification issues with using a window of ages as a proxy for retirement mentioned earlier. However, the small group who continue working through ages 65-74 may not be representative of income and expenditure in the wider population.

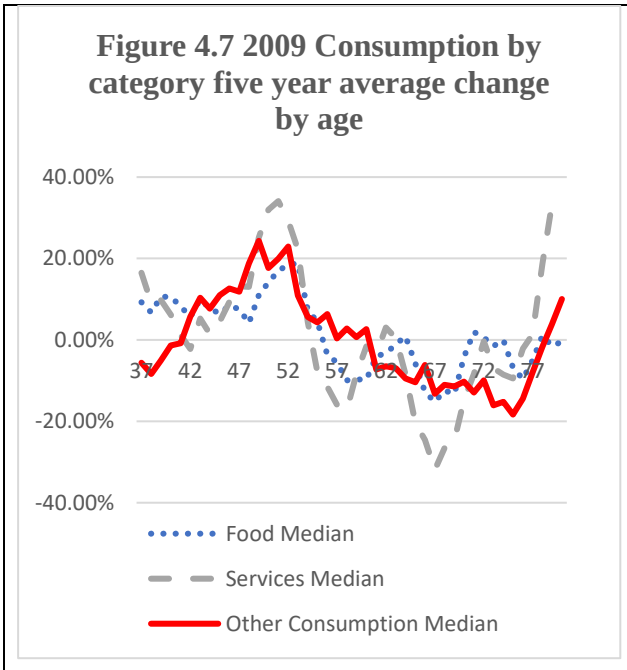

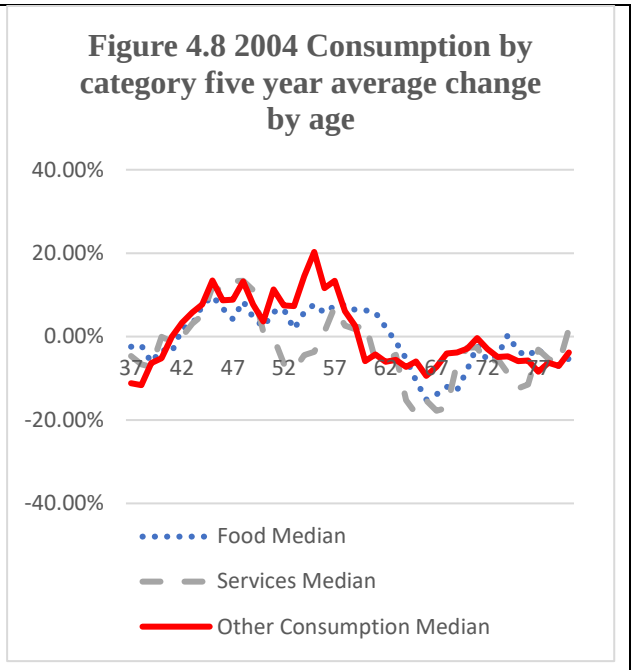

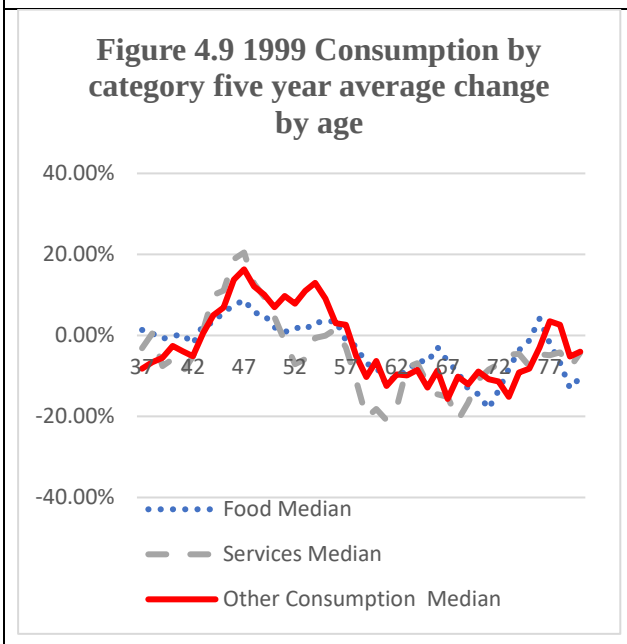

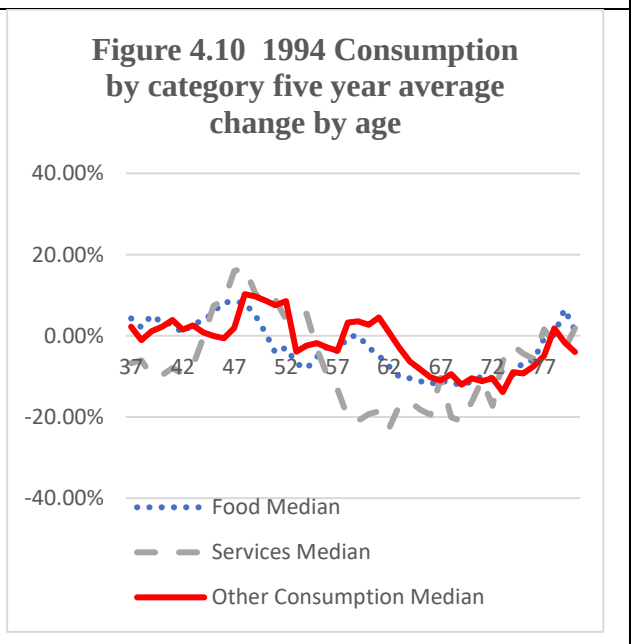

Figures 4.8, 4.9 and 4.10 show similar reductions in food consumption at retirement in 2004, 1999 and 1994 to 2009 of about 15 per cent. The fall in services is more substantial than for food in each year, and particularly so in 1994, which suggests that it is a category that is

particularly sensitive to the retirement transition in Ireland. Other consumption has a similar pattern in all years, showing a slow decline over retirement ages.

4.6 Conclusions and Discussion

Consumption falls at retirement age in the data presented here using cross section data. The results for food show similar results to the UK/US research with a drop of 15-20% at the most common retirement ages between 62 and 67. Otherwise, the main category of expenditure and consumption that falls at retirement age is services expenditure. Falling services consumption in retirement could be a failure of the PIH as it is not considered to be part of work-related expenses in the literature. Services consumption is particularly likely to be exposed to unanticipated income shocks in retirement, since it includes optional expenses like holidays. It is also the category that has recorded the largest increases in working age expenditures in the period 1994-2009. While pensioners have maintained their relative living standards based on incomes during this period, they may be falling behind in these types of discretionary expenditures.

While the results in this section are suggestive of a fall in services consumption in retirement that is inconsistent with the PIH it is important to point out that the analysis relies on cross section data. A full test of the model would require panel data that follows the same individual through retirement, while recording both income and consumption. Unfortunately such data are not available for Ireland.

Appendix 4 Median equivalised income and consumption by age, 1994-2004

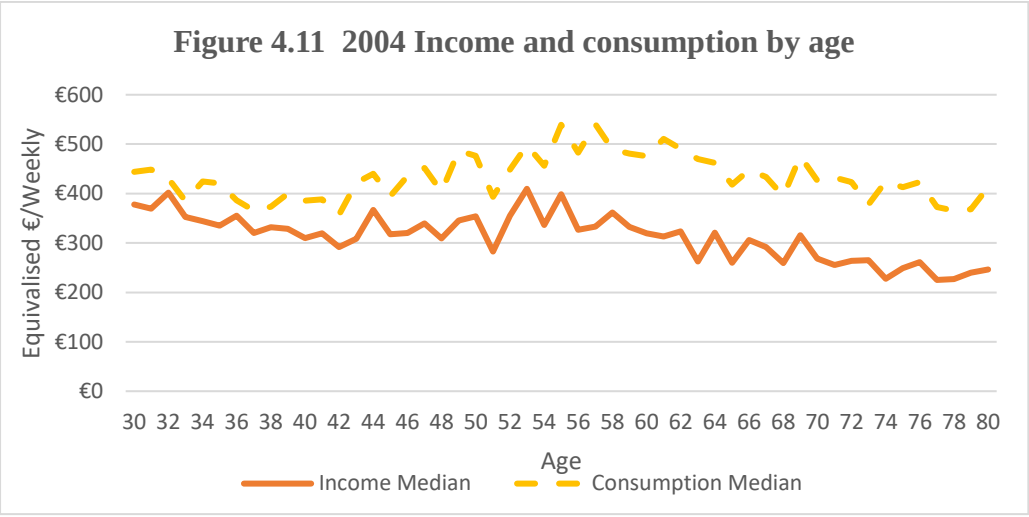

.

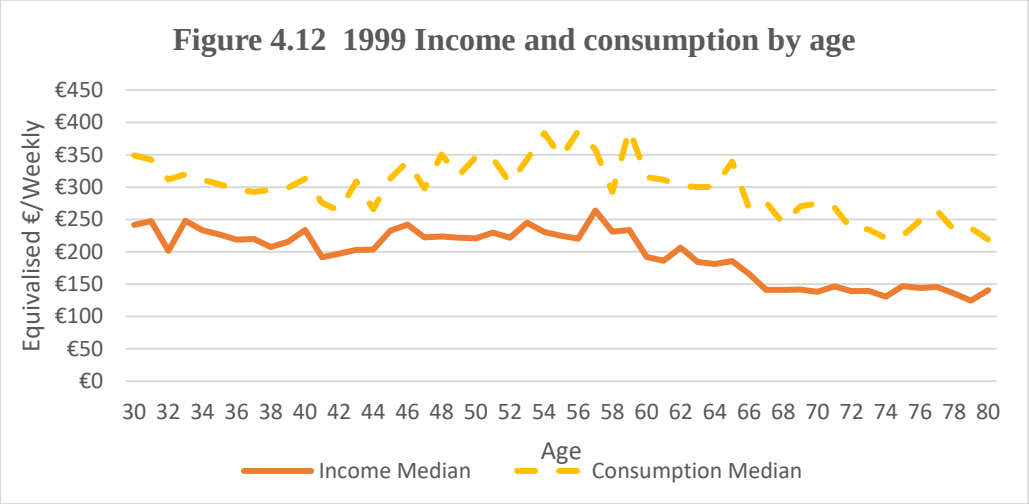

.

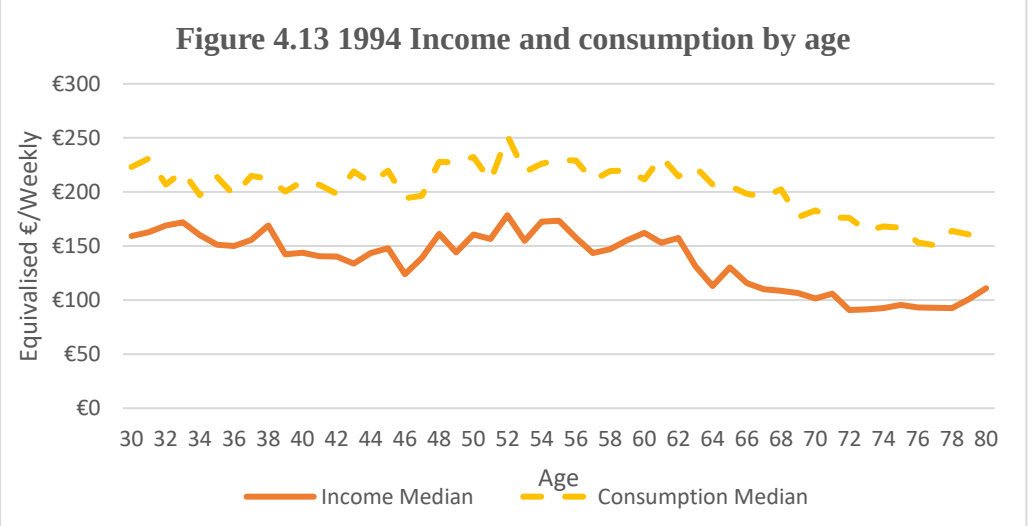

5. Survival expectations in Ireland

5.1 Introduction

Savings by Irish consumers, as defined by income less expenditure in the Household Budget Survey (HBS), are low and concentrated at high income levels. This is inconsistent with the 'Life-Cycle Hypothesis', since consumers are assumed to 'smooth' their lifetime consumption via savings at working age. One possible reason for low savings levels is that consumers underestimate their life expectancy, which is the focus for this chapter. This has implications for pensions policy, since most retirees who have the option choose a drawdown arrangement (or Approved Retirement Fund – ARF[15]) over a guaranteed annuity income, which potentially exposes them to longevity risks later in retirement.

My paper examines this issue by comparing subjective survival probabilities of survey respondents from the Irish Longitudinal Study on Ageing (TILDA), which is a nationally representative study of the population of Ireland aged 50 and above, with objective life expectancy data from life tables, which are calculated using mortality experienced over a three-year period. I do this by firstly testing that survival probability responses are well formed. I use the distribution of responses to check for generic answers, and also use an Ordinary Least Squares regression methodology to establish that responses covary appropriately with risk factors. I then compare the subjective responses to life table data by sex and year of age, which acts as a further check on the validity of the subjective responses. Finally, I examine whether the profile of respondents who appear to misunderstand probability in the survey explains some of the difference between the subjective and objective survival probabilities.

For the first issue on validity of responses, I find that the Irish responses are well formed using internal consistency checks and checks against other expectations questions for respondents who appeared to answer the question inconsistently. The responses also covary appropriately with known risk factors, and both findings are in line with the UK and US research, which is summarised in Section 5.2. Where the Irish results differ is that there is a large block of respondents who appear to have misunderstood the subjective probability question using as they reported 100 per cent survival probabilities, which are of course impossible.

[15] ARFs allow pensioners flexibility in drawing down their retirement incomes compared to annuities, which pay a fixed income. ARFs can also be used for bequest purposes if the pensioner wishes, if funds remain after a minimum pension drawdown of 4 per cent of the fund value per year. Originally ARFs were only available to self employed pensioners. The option was extended to defined contribution scheme employees in 2011.

When I follow the international literature and exclude those reporting 100 per cent survival probabilities from the analysis I find clear evidence of pessimism (i.e., respondents underestimate their actual life expectancy) up to age 70 in the TILDA data compared to the life tables, which is similar to the UK and more recent US research findings. I find that women are more pessimistic than men, also in line with existing research. However, pessimism levels are noticeably higher in the Irish data than in the US and UK. This could be related to lags in information on mortality feeding into subjective expectations, since Irish life expectancy at older ages increased much more rapidly in recent years than in the US and UK.

The main contribution to this research is however on the predictive value of the relatively large group of respondents giving survival probabilities of 100 per cent in the Irish data. When these respondents are included in the comparison with objective data, evidence of pessimism up to age 70 is eliminated. Furthermore, the apparent difference in pessimism between men and women is corrected by including these respondents. Inclusion of this group also corrects an overly smooth profile in subjective responses by year of age, which is unrealistic since the target survival age for respondents changes every 5 years from age 65. The socio-economic profile of this group is similar to those who responded 'accurately', and they also gave a variety of answers to similar expectations questions. These findings suggest that 100 per cent replies are well considered, and that apparent low savings levels in the HBS data are not explained by pessimism. Also, this analysis does not raise concerns about the increased take-up levels of ARF arrangements by pension scheme retirees. Optimism bias for respondents aged 70+ could counteract improvements in life expectancy during retirement, since they are unlikely to draw down their flexible pensions too quickly if they overestimate their remaining life years.

This paper is structured as follows. Section 5.2 provides a literature review. Section 5.3 describes the data used for the analysis. Section 5.4 looks at the validity of responses, with Section 5.5 comparing responses to objective data and Section 5.6 examining respondents who appear to misunderstand probability in further detail. Conclusions are summarised in Section 5.7.

5.2 Literature review

This first study relevant to this topic was Hamermesh (1985), who found that individual estimates of survival probabilities are coherent, useful for prediction and conform to actuarial data in a small sample of academics. Hurd and McGarry (1995) were the first researchers to assess survival probabilities from the Household Retirement Survey (HRS), which is a large

representative survey of the US population aged 51 and over. They found that average survival probabilities were consistent with survival rates from life tables. They also found that average probabilities varied across population groups in the same way as survival rates, with for example, smokers and those self-reporting poor health reporting lower survival probabilities. Hurd and McGarry (2002) added more evidence in favour of survival probabilities from the second wave of the HRS. They found that individuals with higher subjective probabilities in wave 1 experienced lower mortality between the waves, and also that people who experienced serious illnesses between the waves reduced their survival probabilities. Perozek (2008) found that subjective probabilities perform better than life tables in predicting age and gender specific mortality rates in the US, which is the strongest evidence to date in favour of subjective probabilities.

Comparisons against life tables have produced similar results on overall and sub-group level biases in the countries where subjective survival probability data are available. Research by Banks et al. (2004) from the English Longitudinal Study of Ageing (ELSA) found that, as a group, women underestimate life expectancy (or are pessimistic), whereas Gan et al. (2005) and Hurd and McGarry (2002) found that men tend to overestimate life expectancy. Overall, younger age groups of both genders tend to be pessimistic and older groups tend to be optimistic (e.g., Hamermesh, 1985, Elder, 2013, O'Dea and Sturrock, 2018). The issue of how to treat focal point responses for survival probabilities (mainly 50 per cent and 100 per cent) is widely discussed in the literature, as these responses are common across countries and domains. O'Dea and Sturrock (2018) find that respondents who give 50 per cent as their survival probability generally give different answers to similar expectations questions and thus accept that these 50 per cent answers are valid. However, they reject 100 per cent answers on the basis that these respondents do not understand the concept of probability. Hurd (2009) found that zero and 100 per cent answers both had some predictive power using an analysis of 14 year survival rates since 1992, despite being impossible in practice. Gan et al. (2005) developed a Bayesian model using actual survival information to make these focal point answers useable in economic research.

Various approaches have been developed to produce individual level 'survival curves' from subjective probabilities for use in microeconomic models. Gan et al. (2005) developed a model for individual subjective survival curves using a scaling factor called an 'optimism index' which takes respondents' mortality-relevant characteristics into account. Bissonette et al. (2012) modified this approach using actual mortality from the HRS over a 16 year period to

benchmark the survival curves rather than life table data[16], which facilitates analysis of more detailed subgroups. Groups found to be overly optimistic include men, younger groups, black people, the well-educated and smokers. These approaches use a single future age for the survival probability. The TILDA survey also has a single target age per respondent, whereas in the HRS and ELSA, younger respondents are generally asked about survival to two future ages. Since respondents tend to be pessimistic at younger ages and optimistic later, a single target age approach may produce biased results (Elder, 2013). Elder (2013) also found that Perozek's (2008) favourable analysis of subjective probabilities does not hold in general or for population sub-groups due to this bias.

Where additional information is available on target ages, researchers attempt to measure a range of survival expectations (e.g. Wu et al., 2014) or use two data points in survival curves (O'Dea and Sturrock, 2018). Both approaches facilitate more precise estimates of subjective survival curves but find similar results to the earlier analysis. For example, O'Dea and Sturrock (2018) found growing pessimism relative to life tables up to ages in the mid-70's and optimism later in life, and identify particularly high pessimism among women. Both studies relate their findings on pessimism to poor demand for annuities in Australia and the UK, which is also the main policy finding in the US literature. For other pensions issues, Hurd et al. (2004) found that subjective survival only weakly predicted early claiming of social security benefits in the US, while O'Dea and Sturrock (2018) relate poor take-up of pension deferral in the UK to survival pessimism.

The paper most closely related to the present chapter is Nivakoski (2020) who uses the TILDA data to examine whether longer life expectancy leads to larger accumulation of pre-retirement wealth. In contrast, the motivation for my paper is to examine whether consumers underestimate their life expectancy, as found in much of the international literature. I also extend Nivakoski's work by considering the extent to which expectations are well focussed, which is essential for me in setting up my main research question on whether consumers underestimate their life expectancy. In addition, this chapter provides a detailed analysis of the large number of respondents who report a probability of survival of 100 per cent. Individuals answering in this way have been eliminated from previous analyses (e.g. O'Dea and Sturrock) but my analysis shows that inclusion of the 100 per cent responses eliminate apparent

[16] The authors also assume the subjective scaling factor follows a specific distribution - the Gamma distribution instead of the index approach used by Gan et al. (2005)

pessimism levels for respondents aged 50-64, and also significantly reduces gender differences. Inclusion of this group appears to improve the performance of the subjective assessments.

5.3 The data

TILDA is a nationally representative panel study of the population of Ireland aged 50 and above. TILDA aims to understand the health, social and financial circumstances of the older Irish population and how these factors interact to influence the aging process. The first wave of data collection was conducted between October 2009 and July 2011. In total, 8,175 individuals aged 50 and over from 6,279 households participated in the study. In addition, 329 interviews were also conducted with younger spouses or partners of participants, leading to a total sample size of 8,504. I have used the TILDA Anonymised Microdata File from the Irish Social Sciences Data Archive, which top and bottom codes respondents aged below 50 and 80+, so data on these respondents cannot be used since year of age is necessary for much of the analysis in this paper. The usable sample from this file is 7,537 within the age cut-offs as 329 younger and 626 older respondents are excluded and there are also 12 respondents with missing ages. There is limited information loss from focussing on this age range since expectations at higher ages are found to be unrealistic in later analysis.

The subjective survival expectations question was only asked in wave 1 of TILDA, and is similar to the UK and US panel study questions referenced above:

IWER: SHOW CARD EX1
EX101. Using the scale on this card, what is the percent chance that you will live to be
IWER: CODE PERCENTAGE GIVEN BY RESPONDENT
75 IF AGE IS LESS THAN 65
80 IF AGE IS 65-69
85 IF AGE IS 70-74
90 IF AGE IS 75-79
95 IF AGE IS 80-84
100 IF AGE IS 85-89
110 IF AGE IS 90 OR OVER
0...100
98. DK
99. RF
(HRS)
BL: IF Rs' AGE <62 GO TO EX102
BL: IF (Rs' AGE >=62 AND <65) GO TO EX103
BL: IF Rs' AGE >=65 GO TO EX104

Where the scale presented to the respondent is from 0 to 100;

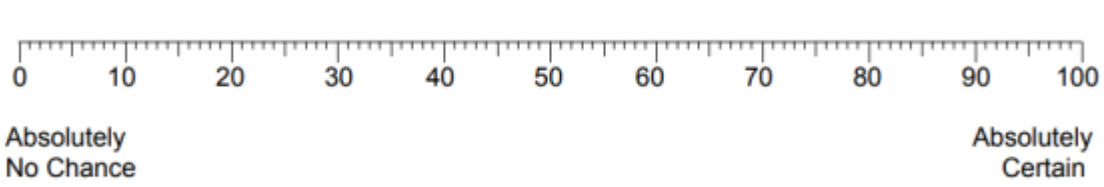

A wide range of household and individual level data are collected in the survey, which offers the possibility of controlling for background factors in relation to health, education, household structure and many other dimensions in economic analysis.

Results from the CSO Irish Life Tables are compared against subjective probabilities in Section 5. Life Tables were constructed for males and females which are representative of the mortality experience in Ireland in 2011 by using the 2010 and 2012 population estimates and 2011 Census of Population (usually resident) combined with deaths registered in the three years. The life table should reflect the normal mortality conditions at about the time of the Census. The life tables report period 'point in time' life expectancy estimates, which are not strictly comparable with TILDA cohort data. Figure 5.1 shows that the trend in life expectancy at age 65 has been strongly positive since 1996. If this trend continues into the future, period life expectancies will understate the true life expectancy of the current generation of older people, since they only reflect the current mortality experience.

Figure 5.1 Period life expectancy at age 65, 1970-2012

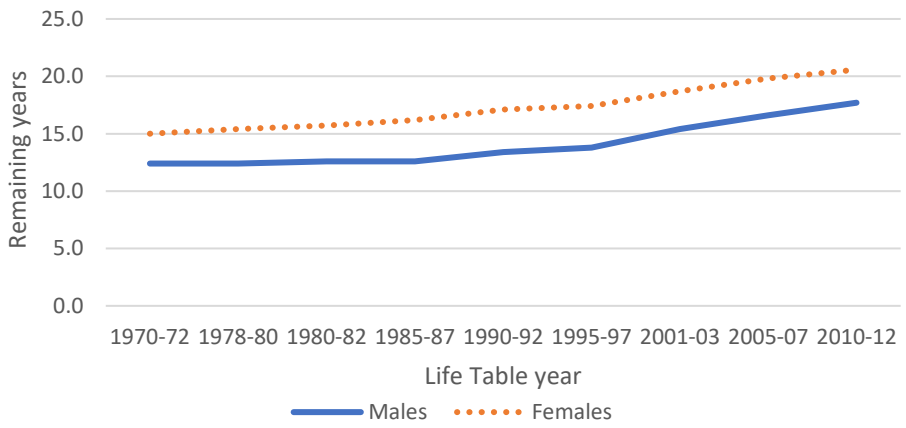

In the UK, cohort life tables are available from the Office for National Statistics in addition to period life tables. The difference between UK cohort and period life expectancies for 2011 by gender is used in the analysis in Section 5.5 as a cohort adjustment factor for each year of age in the Irish Life Table. Cohort life tables are calculated using age-specific mortality rates,

which allow for known or projected changes in mortality in later years. A cohort life table provides mortality rates that vary over time for each age. For example, cohort life expectancy at age 65 in 2014 is worked out using the mortality rate for age 65 in 2014, for age 66 in 2015 for age 67 in 2016 and so on. The cohort estimates use observed mortality rates in 2014 and projected mortality rates from 2015. Therefore, cohort estimates are regarded as a more appropriate measure of how long a person of a given age is expected to live on average than period life expectancy. High and low variant mortality projections are provided in the UK cohort tables. I have used an average of the two as an adjustment factor, which is conservative compared to the recent Irish experience of increasing life expectancy described in Figure 5.1.

5.4 Are the responses well formed?

In this section, I will assess whether responses to the subjective expectations questions are well formed using the distribution of probabilities reported and answers to other expectations questions in the TILDA survey. This will be the basis for more detailed examination of the responses in terms of correlations with health outcomes and comparisons against objective outcomes.

In the US and UK research, respondents who state a 100 per cent probability of survival are generally taken to have misunderstood the question since they do not grasp the concept of probability. There were 2,316 respondents in this category in TILDA, giving a 31 per cent rate of misunderstanding for the question. This compares to 7 per cent and 15 per cent in the UK and US equivalent surveys respectively. In addition to this issue, responses to the question clustered around focal points, particularly at the 50 per cent level but also at 80 and 90 per cent. This can be seen in Figure 5.2, which shows the distribution of survival probabilities excluding 100 per cent responses. The three highest frequency focal point responses of 50, 80 and 90 per cent account for around half of responses in the range 0-99. Clustering was also found in the UK and US surveys and was taken as further possible evidence that respondents did not understand the question, since it could be construed as a generic response.

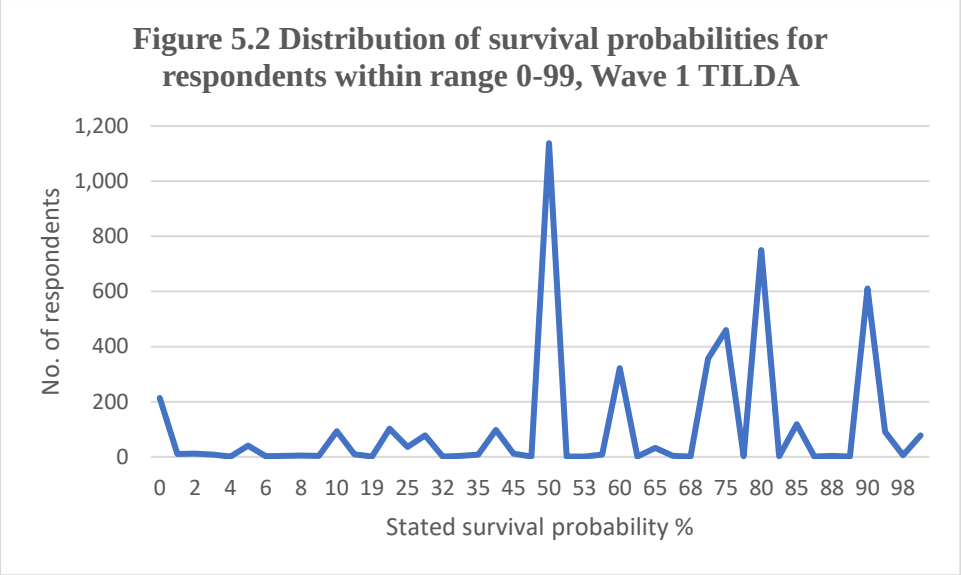

TILDA also included three additional expectations questions that were addressed to all of the respondents who were asked the subjective survival question. These questions were subjective probabilities relating to entering a nursing home, independent living and memory difficulties, and are included in Appendix 5. If respondents also misunderstood these questions, it was taken as a further sign that subjective survival expectations were not well considered in the research. Figure 5.3 presents the percentage of respondents who also gave 100 per cent or 50 per cent responses to the other subjective probability questions. This shows that there was a high level of overlap between those who gave 100 per cent responses and less of an overlap for 50 per cent responses. The 50 per cent responses can be regarded as valid on this basis. Comparable results were found in the ELSA survey by O'Dea and Sturrock (2018). While the overlap for 100 per cent responses is higher, there is a very low proportion of respondents who gave 100 per cent responses to all the questions which does not suggest a complete misunderstanding of probability by respondents

Figure 5.3 Percentage of respondents who misunderstood subjective survival probability and similar questions

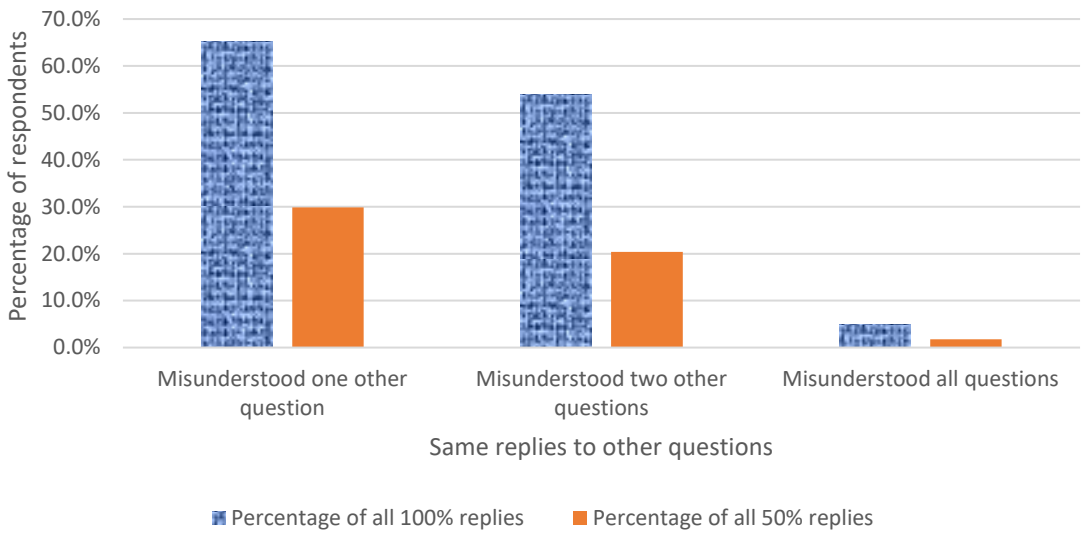

The main distinguishing factor for the 100 per cent group is that they are concentrated in younger ages. Figure 5.4 shows the age distribution of 100 per cent responses. A distinct level shift in responses can be seen at ages 64-66 from 35-40 per cent to below 25 per cent, with the average across all ages at 31 per cent. The profiles of the 100 per cent group by socio-economic status and education level are very similar to other respondents. Using a 5 percentage point difference threshold (i.e. 5 per cent above or below the overall average share of 100 per cent responses), women and those in 'excellent' health had significantly higher shares than average (36 and 46 per cent), while people who experienced heart related or cancer episodes had significantly lower shares of 100 per cent responses (19 and 25 per cent). This suggests that there is a factual basis for a positive assessment of survival prospects by the 100 per cent group. Table 5.1 compares the composition of the 100 per cent group to those with high survival probabilities (90-99) and the rest of the sample. Apart from the larger share of female respondents, the group appear to more in line with high probability group than other respondents. On this basis, 100 per cent responses might be regarded as approximately valid. Section 5.6 addresses whether these responses improve the comparison of subjective probabilities with life table outcomes by age and sex.

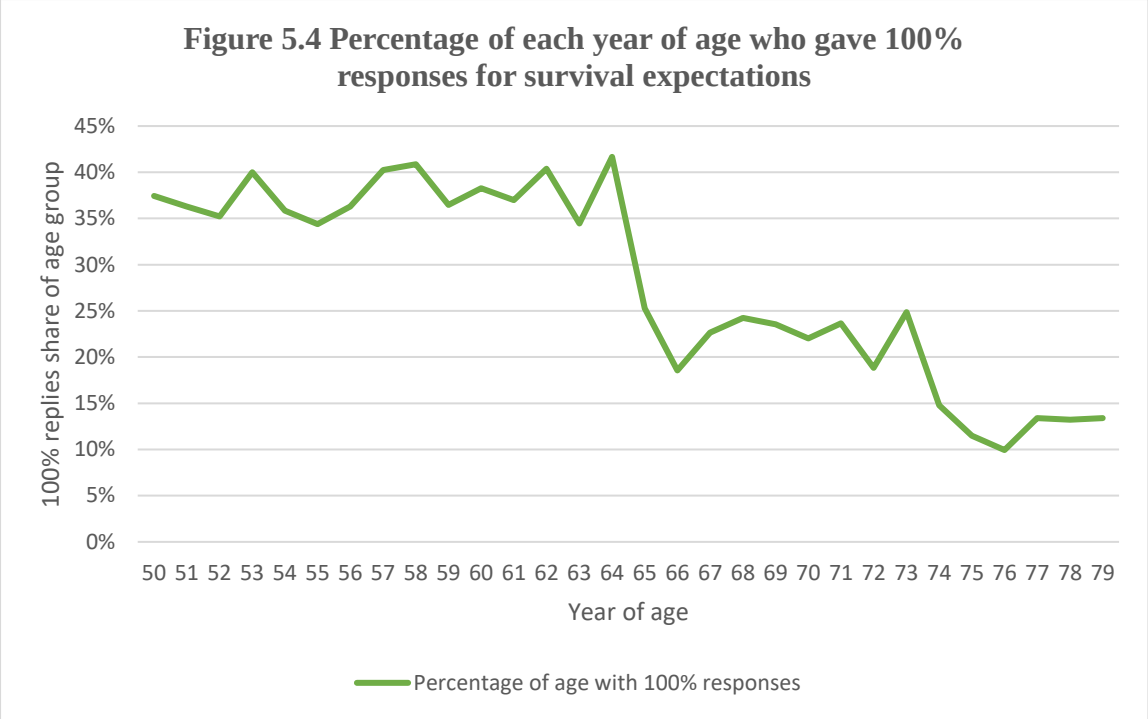

Table 5.1 Comparison of 0-89, 90-99 and 100 per cent groups by characteristic

Characteristic	Age/share of 90-99% group	Age/share of 100% group	Age/share of 0-89% group
Average age	61.0	60.1	63.2
Share of females	53.8%	63.1%	48.0%
Share of males	46.2%	36.9%	52.0%
Share with heart attack	2.3%	2.6%	5.4%
Share with cancer diagnosis	3.1%	5.1%	7.5%
Share in ‘good’ or 'excellent' health	85.6%	85.6%	72.0%

My next step is to look at the risk factors and outcomes that are associated with having higher or lower expectations of survival to older ages. This analysis will assess whether responses reflect well-formed judgements about future mortality. If the expectations correlate with outcomes that are relevant to mortality, this can be taken as evidence that they are well formed and likely to drive economic decision making.

Table 5.1 shows the results of an Ordinary Least Squares regression analysis looking at the relationship between these risk factors and subjective survival probabilities for all reported probabilities between 0 and 99 per cent. The reported numbers for the included risk factors represent the mean percentage point difference in the dependent variable of subjective

probability for someone with the risk factor compared with someone without it, holding the other risk factors constant. The analysis is run separately for men and women to show the difference in impacts by sex. Also, self-reported health is included as a risk factor in the third and fourth columns to identify whether the risk factors have an additional impact for respondents otherwise in good health. While not all of the results are statistically significant, most of the relationships identified are similar to those found in previous research. For example, smoking (previous and current) is a clear risk factor as are previous serious health conditions. The positive relationship between alcohol and survival probabilities is surprising, and there are three aspects to this:

- Research for other countries (O'Dea and Sturrock for the UK, Hurd and McGarry 1995 for the US) finds positive relationships between moderate alcohol consumption and survival probabilities in most cases;
- There are no controls for age or other background factors in this initial analysis, and alcohol consumption is higher in younger age groups who generally experience better health;
- As the survival expectations question uses a different target age for different age groups, this also needs to be controlled for.

Table 5.2 OLS Regression: The Impact of Risk Factors on Subjective Survival Probability

Risk factor	Excl. self-reported health		Incl. self-reported health	
	Male Coefficient	Female Coefficient	Male Coefficient	Female Coefficient
Smoking (relative to non-smoker)				
Ex-smoker	-3.5***	-3.0**	-2.7**	-2.0
	(1.236)	(1.341)	(1.209)	(1.302)
Current smoker	-8.4***	-7.5***	-6.4***	-5.6***
	(1.637)	(1.544)	(1.610)	(1.505)
Alcohol consumption (relative to 0-3 drinks)				
4-10 drinks per week	5.1***	4.5***	3.8***	3.0**
	(1.386)	(1.491)	(1.359)	(1.450)
More than 10 drinks per week	0.8	5.3***	-0.4	3.8**
	(1.313)	(1.878)	(1.286)	(1.822)
Health conditions				
Heart attack	-14.7***	-9.9**	-10.8***	-3.9
	(2.161)	(3,975)	(2.145)	(3.888)
Cancer diagnosis	-9.8***	-8.6***	-7.2***	-6.4***
	(2.311)	(2.150)	(2.269)	(2.091)

Statistical significance at the 10%/5%/1% level is denoted by */**/***, Standard Errors in parentheses.

Table 5.2 reports results from the same regression model as Table 5.1 with controls for age and education added. The age control combines single years into five-year age categories and so also controls for the change in target age in the survival expectations question. The correlations identified by the coefficients are generally in the direction that would be expected from other evidence about the impact of these risk factors on life expectancy. The relationship between alcohol consumption and life expectancy has low correlations, and this was also found in the UK research (O'Dea and Sturrock). There are generally significant negative correlations between previous health conditions and survival probabilities, which is evidence in favour of the responses being well formed and reflective of the life experience of respondents. When 100 per cent responses are included, all the signs of the coefficients remain the same and are at most one standard error lower or higher than reported in Tables 5.2 and 5.3, so inclusion of this group does not affect the conclusions from the restricted sample.

Table 5.3 OLS Regression: The Impact of Risk Factors on Subjective Survival Probability including Age/Education controls

Risk factor	Excl. self-reported health		Incl. self-reported health	
	Male Coefficient	Female Coefficient	Male Coefficient	Female Coefficient
Smoking (relative to non-smoker)				
Ex-smoker	-2.3*	-3.0**	-1.7	-2.2*
	(1.217)	(1,293)	(1.186)	(1.262)
Current smoker	-8.3***	-9.2***	-6.6***	-7.7***
	(1.622)	(1.522)	(1.587)	(1.489)
Alcohol consumption (relative to 0-3 drinks)				
4-10 drinks per week	2.9**	1.0	0.7	-0.04
	(1.387)	(1.470)	(1.355)	(1.434)
More than 10 drinks per week	-0.4	1.6	-1.3	1.0
	(1.309)	(1.839)	(1.752)	(1.792)
Health conditions				
Heart attack	-12.1***	-5.2	-8.0***	-0.6
	(2.131)	(3.851)	(2.112)	(3.777)
Cancer diagnosis	-7.4***	-8.6***	-4.8**	-5.8***
	(2.274)	(2.074)	(2.227)	(2.026)

Statistical significance at the 10%/5%/1% level is denoted by */**/***, Standard Errors are in parentheses.

In summary, subjective expectations seem to be well formed for TILDA respondents.

5.5 How do responses compare to published life tables?

In this section, I compare average subjective survival probabilities for men and women against objective survival probabilities from Irish Life Tables. These life tables are compiled from population outcomes and thus reflect true mortality levels experienced in the population.

The TILDA expectations results for probability of survival to future target ages are compared to the Irish Life Tables for males and females below in Figures 5.5 and 5.6. In general, both males and females are pessimistic (on average they overstate future chances of mortality) up to age 70 and then are optimistic from ages 70-79, which is in line with the UK and US literature discussed earlier. The earlier period is more relevant to planning and smoothing outcomes under the Life-Cycle Hypothesis as pensioners are hypothesised to run down savings post retirement. The TILDA average results could be expected to display 'steps' given that a different question is asked of every five-year age group from age 65. In fact, TILDA expectations are quite smooth compared to the life table estimates where the target age for

calculating survival rates is also shifted every five years to match the TILDA question. I have included a cohort adjustment from the UK tables as discussed in Section 5.3 for Irish period life expectancies and males and females are even more pessimistic up to age 70 when this adjustment is included.

The male and female patterns are similar, but Figure 5.6 shows that females are relatively more pessimistic than males (18 p.p. vs 12 p.p. average difference from ages 50-64 with the cohort adjustment), which is also found in the research. This comparison suggests that women are possibly not well informed about their favourable life expectancy outcomes. There was a similar gap between females and males in the UK research by O'Dea and Sturrock (2018) with average differences of 10 and 6 p.p. for all born in 1940-49 using a target survival age of 75. This also suggests a similar information asymmetry in the UK.

Figure 5.5 Male subjective survival probabilities compared to population estimates by year of age

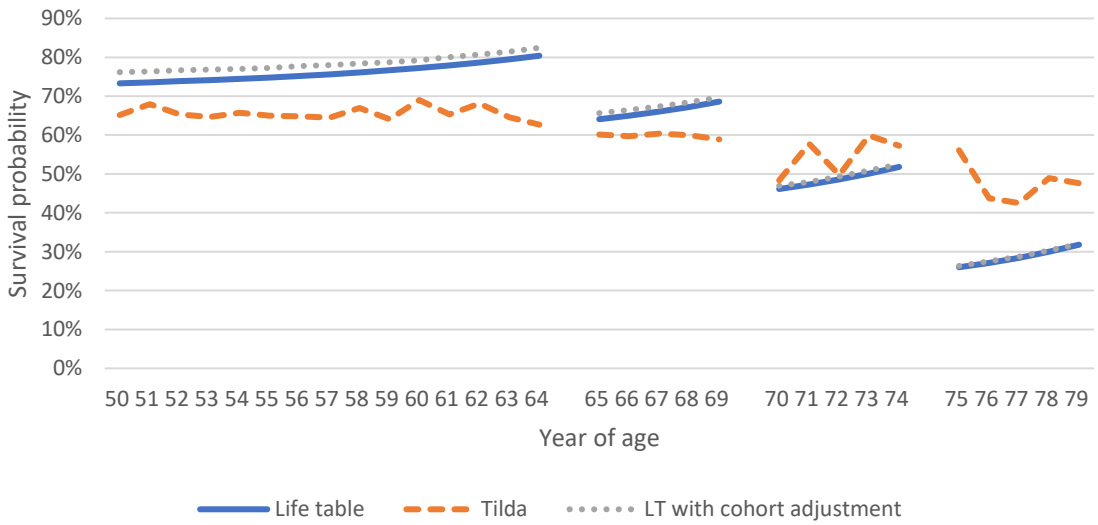

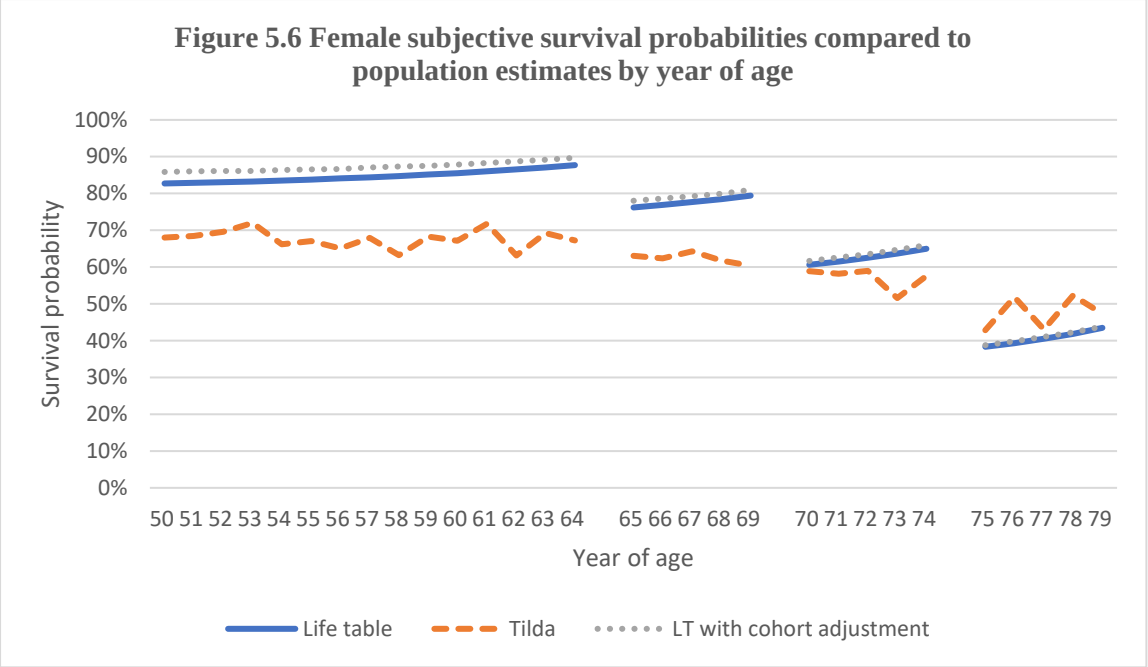

5.6 Do respondents who misunderstand probability affect the conclusions on pessimism?

In this section, I repeat the analysis from Section 5.5 including the relatively large group who gave 100 per cent responses for subjective survival probability. This will assess whether the findings on pessimism are still valid on this basis.

Figure 5.7 presents the male comparison of survival probabilities including the 100 per cent group with period life table estimates. I have excluded the cohort adjustment line as it makes little difference to the comparison. Also, while this adjustment is relevant to the accuracy of survival forecasts, future mortality does not necessarily contribute to understanding of the question. Respondents are more likely to reflect on current than future mortality experience when answering the question. Figure 5.7 shows that pessimism is eliminated up to age 70 when this group are included, while optimism bias from age 70 onwards is further reinforced.

Figure 5.7 Male subjective survival probabilities compared to population estimates by year of age incl. 100% responses

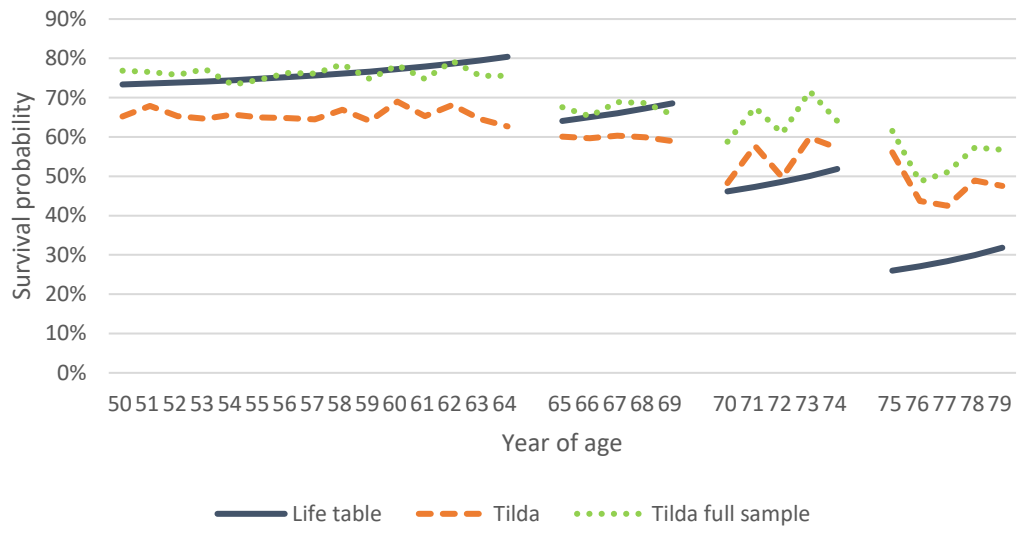

Figure 5.8 Female subjective survival probabilities compared to population estimates by year of age incl. 100% responses

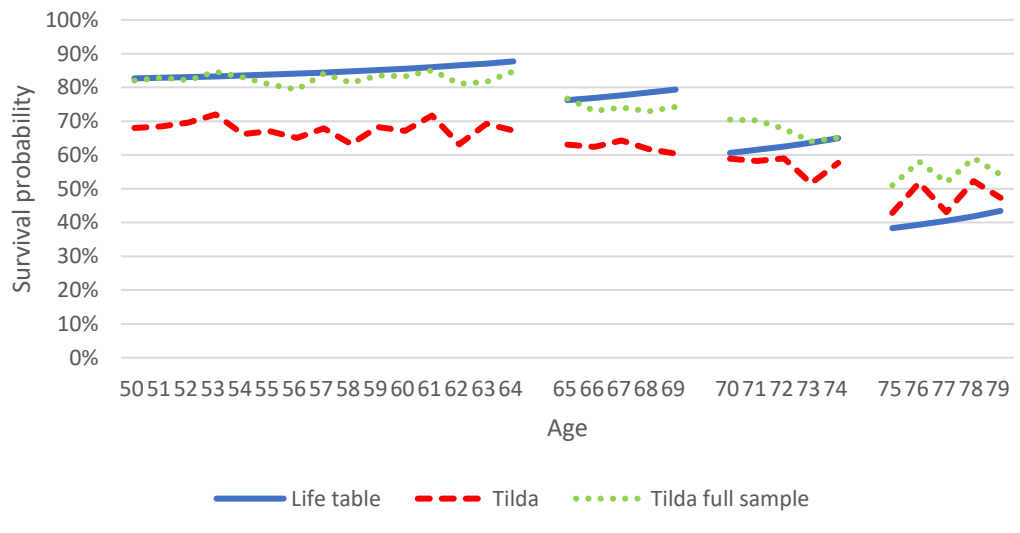

Figure 5.8 shows the female comparison, which gives further insight into the predictive value of these 100 per cent respondents. The high pessimism bias for women found in Section 5.5 is eliminated by their inclusion, as women are overrepresented in this group. Finally, the overly smooth profile of respondents found in Section 5.5 is replaced with steps when the target age changes in the question. This suggests that the responses are well distributed by respondent year of age within the different target age groups. These two additional findings particularly suggest that there is strong predictive power in these responses, since apparent inconsistencies or biases in the Section 5.5 profile are corrected when they are included.

5.7 Conclusions and Discussion

I compare subjective survival probabilities from TILDA with actual survival rates from life tables in this analysis. Based on the criteria used in most of the international research, I find evidence of survival pessimism up to age 70 and optimism afterwards, as well as additional pessimism bias for women. These findings are in line with the research. However, there is a large group of TILDA respondents who gave 100 per cent replies for their survival probabilities. This group is generally either excluded or corrected for in the international research. When this group is included in my analysis, pessimism bias below age 70 is much reduced and optimism bias from 70+ is reinforced. Inclusion of the group also 'corrects' additional pessimism bias for women and inconsistencies in responses at different target ages. I find that these 100 per cent responses are approximately valid in the TILDA survey and have predictive value for future survival prospects. This highlights the importance of carefully considering all responses in such an analysis, as well as the potential bias that may arise when using a restricted sample.

My preferred estimates imply that pessimism does not explain apparent low savings levels in the HBS, and also raises no concerns about Approved Retirement Fund members in relation to assessment of longevity risks when drawing down pension incomes.

Appendix 5 Similar expectations questions from TILDA

IWER: SHOW CARD EX1
EX104. Using the same scale, what are the chances that you will move to a nursing home in the next five years?
IWER: CODE PERCENTAGE GIVEN BY RESPONDENT
0…100
98. DK
99. RF
(HRS)
Note: Nursing homes are institutions primarily for people who need constant nursing supervision or are incapable of living independently. Nursing supervision must be provided on a continuous basis for the institution to qualify as a nursing home. Please don't include stays in adult foster care facilities or other short-term stays in a hospital.
BL:
IF Rs' AGE <75 – GO TO EX105
IF Rs' AGE >75 – GO TO EX107

IWER: SHOW CARD EX1
EX105. Using the same scale, and assuming you will still be alive at age 75, what are the chances that your health will allow you to live independently? By that, I mean to live at home without help and to manage your own affairs?
IWER: CODE PERCENTAGE GIVEN BY RESPONDENT
0…100
98. DK
99. RF
(HRS)

IWER: SHOW CARD EX1
EX106. Using the same scale, and assuming you will still be alive at age 75, what are the chances that you will be free of serious memory or reasoning problems? By that I mean problems that would interfere with your ability to manage your own affairs?
IWER: CODE PERCENTAGE GIVEN BY RESPONDENT
0…100
98. DK
99. RF
(HRS)

6 Volatility of earnings and income in Ireland, 2005-2017

6.1 Introduction

Inequality has been widely examined using repeated cross section data, and the impacts on society have been discussed by Piketty (2014) and others as increased inequality has often coincided with negative social outcomes. Deaton and Paxson (1994) related income variance at individual level to inequality outcomes and found that increased permanent variance results in higher levels of aggregate inequality over time. Variance components analysis of incomes has focussed on the nature of inequality for consumers. Income is analysed longitudinally to assess the extent to which cross section inequality represents differences between individuals that are long lasting/permanent or transitory. Higher transitory variance translates into higher income risk since it is not anticipated by consumers. In a PIH model individuals may be expected to reduce consumption in the face of these shocks. However, if individuals can insure themselves against labour market income (earnings) risks arising from unemployment or pay cuts by using welfare, running down savings or income sharing within households this may allow them to consumption smooth in the face of these shocks. Comparisons of the trend in transitory variance based on successively broader definitions of household income make it possible to identify the insurance provided by each additional income source.

Irish analysis of inequality has found that it has been quite stable over the long term, including during the "Great Recession" (Callan, Bercholz and Walsh, 2018). Irish research on inequality has mainly been based on annual cross section Survey of Income and Living Conditions (SILC) data. Volatility analysis requires panel data, which involves repeated observation of the same people or households over time. There has been no Irish analysis of volatility to date due to data availability. I combine individual administrative data on earnings and social welfare from 2005-2017 with data on education and household linkages from the 2015 Household Budget Survey (HBS) to obtain new insights on income volatility during the "Great Recession". This analysis reveals that male working age earnings risk increased significantly during 2008-10 to an extent not identified in Irish inequality analysis so far. The increase in male earnings risk is also very high in an international context, although counter-cyclical volatility increases have also been found by US and UK researchers (Shin and Solon, Jenkins, both 2011). The increase in volatility found in Ireland is concentrated in low income and low education groups. A

combination of job losses, pay cuts and hours reductions are reflected in annual earnings data in the panel, which combine to drive the increase in volatility. A second peak in earnings volatility is identified in 2013-2014 as employment, hours and earnings simultaneously improve rapidly for men, and this post-recession 'spike' is unusual in an international context. The results show that male earnings volatility in Ireland falls back to below pre-crisis levels after 2014, whereas some US researchers have found a longer-term increase in volatility or transitory variance after crises. Earnings volatility remains high in 2015-17 in Ireland by international comparisons despite the fall from peak levels. Availability of administrative social welfare incomes and household linkages from HBS facilitates an analysis of 'insurance' provided by male welfare and spouses' incomes. Both combine to result in much flatter and lower profiles of volatility.

6.2 Literature review

Income inequality is usually examined using cross section survey results, since that is often the only income data available to researchers. There is a significant body of Irish research on inequality based on the annual SILC survey. Most of the research, discussed below, finds that the trend in inequality of Irish incomes has been quite stable over the past 30 years despite the two deep crises that were experienced during that time. Incomes are composed of earnings of household members plus social welfare less taxes, so while overall inequality has been steady, the components have been more variable. There have been periods where social welfare payments caught up with earnings increases over time or the tax system changed to reduce inequality in market incomes.

Callan, Bercholz and Walsh (2018) found that, despite recessions, household incomes have grown strongly in Ireland in the last 30 years with stability across the income distribution. Ireland's rapid and equal growth in incomes across the distribution is unusual in an international setting. During this time, income inequality has risen in many other developed countries, which has resulted in Ireland moving from a high-ranking country among this group to a middle ranking. Market income inequality is high in Ireland, but a redistributive tax and welfare system has helped to offset that. Over the 1987 to 2014 period, discretionary changes in tax and welfare policy led to gains that were greatest among those with incomes in the lowest 20 per cent of households. Callan, Doorley and Savage (2018) examined the crisis period of 2008-2013 in more detail using the SWITCH tax/benefit model for Ireland and similar microsimulation models for other countries who experienced a deep recession. They found that

stability of inequality during this period was due to the 'automatic stabiliser' of increased welfare take-up rather than changes in market incomes or tax/benefit policies. Roantree (2020) examined income inequality using a broader range of data than Callan, Bercholz and Walsh over the same 30-year period and found that inequality at the lower end of market incomes decreased due to improved hourly wages. He also found that female participation improvements during the period fed into more unequal higher incomes, since many of the women who benefitted were married to higher earners. Madden (2014) and later Dooley and Madden (2022) found that growth was broadly pro-poor over three periods (2003-2007, 2008-2011 and 2012-2019) based on growth incidence curves.

Earnings inequality during the Irish crisis has also been examined, including an analysis of panel data. Doris et al. (2015) used administrative longitudinal data to follow individual earnings for the entire employee population in Ireland between the years of 2005 and 2013, which corresponds to the first half of the period in my analysis. They found that the Irish labour market was flexible in the pre and post crisis periods. The proportion of workers receiving earnings cuts more than trebled during the crisis. In addition, these wage cuts were progressive, particularly in the public sector, where the highest wage earners recorded cuts to earnings of 12 per cent. However, their analysis was limited by the lack of control variables for individual characteristics over time in the panel data, which made it difficult to analyse changes in inequality. Holton and O'Neill (2017) extended the analysis of wage inequality in Ireland by examining hourly wage dispersion from 2004 to 2013 using cross section SILC data. A detailed comparison of their findings on earnings with the panel used in this paper is included in Section 6.3. They decomposed changes in wage inequality into a component due to changes in the price of skill and a component due to changing characteristics of the workforce and found lower returns to skills during this period despite increasing education levels.

Cross section analysis facilitates measurement of the level of inequality in a year, and a comparison of trends over time. Economists are also interested in longitudinal analysis of inequality, with the main purpose of identifying whether changes to income are permanent or transitory. Higher transitory variance translates into higher income risk since it is not anticipated by consumers. Panel data, where the same person or household is observed repeatedly over time, is required for this type of variance components analysis. Much of the research on variance components analysis is from the US, Canada and UK, where these data sources are available.

Moffitt and Gottschalk (2002) presented both 'simple' and econometric based models of variance components analysis based on US panel survey data. The simple model is discussed in detail later, while the econometric model uses minimum distance estimation for the permanent and transitory components (Gottschalk and Moffitt, 1994). They found that the variance of permanent male earnings began rising in the late 1970's and continued rising through the 1980's, so that permanent inequality was the main driver of increased inequality over the period. The variance of transitory earnings also rose in the 1980's but declined in the 1990's in the results of the 2002 analysis; subsequently Moffitt and Gottschalk (2012) found that the transitory variance remained at the high level of the late 1980's through to 2004. This aspect of their research was particularly influential, as increased unanticipated income volatility or instability makes household consumption and retirement planning more difficult.

Shin and Solon (2011) raised issues with both the simple and more formal approaches of Moffitt and Gottschalk and had particular difficulty with the transitory component being used as a measure of income volatility. They said that variance components models are 'arbitrary mechanical constructs'. To illustrate this, they pointed to Baker and Solon's (2003) rejection of Moffitt and Gottschalk's restrictions based on more detailed Canadian data and different US findings by Haider (2001) based on a change in model specification. Meghir and Pistaferri (2010) presented an overview of the wide range of specifications used for variance components analysis by researchers. Shin and Solon (2011) also pointed out that conclusions among researchers who use the Panel Survey of Income Dynamics (including Moffitt and Gottschalk, Haider and Shin and Solon) diverge in their opinions about the timing of the permanent and transitory shifts. Their proposal was to use an earnings volatility measure based on the dispersion in year-to-year earnings changes reflecting both permanent and transitory shocks, which in their view is much more relevant to volatility measurement than variance components models. Unlike Moffitt and Gottschalk, they found that the increase in men's income volatility occurred in the 1970's with no clear subsequent trend until an increase after 1998. While most of the US volatility studies find that the variance of income shocks is countercyclical (i.e., lower during periods of growth), Guvenen, Ozkan and Song (2015) found that left skewness of shocks is counter-cyclical and shocks are not counter-cyclical on average, i.e, large upward earnings movements are less likely during recessions and large drops are more likely. They also found that earnings for the top 1 per cent are considerably more pro-cyclical than the rest of the population.

Jenkins (2011) applied the US methodologies to UK panel data from 1991 to 2006, including a complete analysis of volatility in household income components. His main findings were that the increase in transitory income volatility found in all the US studies is not clear in the UK data. Unemployment fell in the UK throughout the period of analysis, so the same pattern of counter cyclical volatility is observed in the UK. He found evidence of offsetting effects in household income components, with volatility increasing for spouses' labour income throughout the period and for benefits from 1999 on. Stability in other household income components offset these increases in average incomes.

6.3 Data

Previous work in Ireland has been unable to properly analyse income volatility due to the lack of suitable panel data. I overcome this barrier by constructing a new comprehensive panel data of income by linking administrative data from 2005-2017 with data from the 2015 HBS. Linking the administrative data to the HBS allows me to include covariates not typically available in administrative panel data. My income panel covers 13 years of administrative income data (2005-2017) for most 2015 Household Budget Survey (HBS) respondents. Linkage is possible for almost 90% of HBS respondents and is almost complete for working age groups. The earnings data are originally from income tax sources and covers all income liable for social insurance, which includes self-employed income, occupational pensions and additional self-assessed[17] taxable income such as rent, investment income and shares/share options. Earnings are recorded on an annual basis in the administrative dataset I am using, which is the Department of Social Protection's Central Records System. Social insurance credits are also available in the dataset and are recorded for all unemployment and illness/invalidity welfare claimants. While they are not a comprehensive measure of welfare receipt, groups who experience an income shock generally receive credits. I used weekly welfare credits (1-52 per year) multiplied by weekly welfare rates relevant to each year to estimate annual social welfare incomes. Some welfare benefits are taxable, which may result in double counting between earnings and welfare in gross incomes based on this approach. This is not likely to be an issue for workers who have suffered significant income losses due to long term unemployment or illness as they are usually not liable for income tax.

In addition to the administrative data, personal and household characteristics from the HBS are also available. The main added value from the HBS is that it identifies the household

[17] The income base for social insurance was expanded to include these incomes in 2012.

relationships between individuals in the panel. Household income volatility can then be compared to individual earnings and income volatility. I have only used earnings and estimated welfare incomes for the first and second earner in each household for this analysis. Individual age, sex and education levels are the main HBS variables that I have used in addition to the household linkages in the panel. Grossing factors are also available in the HBS but are on a full population basis and are only relevant to 2015 so I have not used them in this analysis.

I also used age and sex from the 2015 HBS to select the sample for the panel. The final sample is unbalanced, since each year includes different individuals depending on whether they have reported incomes and meet the age criteria in each specific year. All males aged 22-70 in HBS 2015 who were first or second earners in their household were included in the initial sample before age restrictions were applied. This gives a base sample of 3,517 individuals. The panel sample for each year was then selected based on two criteria. The first was that only reported incomes in each year were included in the sample. All returns to Revenue and estimated welfare incomes were used apart from zero incomes and the top 1 per cent in line with international research (Shin and Solon, 2011[18]). The second step was to restrict the sample to ages 25-60 in each year of the panel. For example, the 2005 sample included individuals who were aged 35-70 in HBS 2015, as illustrated below:

Earnings year	Low age current year	High age current year	Low age 2015	High age 2015
2005	25	60	35	70
2006	25	60	34	69
2007	25	60	33	68
2008	25	60	32	67
2009	25	60	31	66
2010	25	60	30	65
2011	25	60	29	64
2012	25	60	28	63
2013	25	60	27	62
2014	25	60	26	61
2015	25	60	25	60
2016	25	60	24	59
2017	25	60	23	58

[18] Shin and Solon also excluded the bottom 1% of earnings, but in this analysis it would randomly exclude individuals in a group with similar levels of social welfare benefits.

The resulting panel dataset is in wide format with six variables for each year (age, earnings, welfare. male gross income, partner's gross income and household gross income) with blank entries for individuals in the base sample who do not meet the income and age criteria. Male education levels in 2015 are also available in the panel.

Table 6.1 shows summary statistics for male earnings over the period. Male earnings growth has been higher at higher earnings levels during this period. Male earnings at the 25th percentile were 13 per cent higher in 2017 than in 2005, while male earnings at the median and 75th/95th percentile levels were 20 and 34 per cent higher respectively over the period. Table 6.2 presents corresponding statistics for incomes, which include social welfare, and shows a similar pattern of increasing inequality over the period. Male incomes were 11 per cent higher at the 25th percentile level, 19 per cent higher at the median and 33-34 per cent higher at the 75th and 95th percentile levels. Average incomes are lower than average earnings in all years since they are calculated over a larger sample including those with only social welfare incomes, which are generally lower than earnings. Standard deviations of earnings and income increase by around 50 per cent over the period. Sampling errors grow in line with income and earnings and do not affect any conclusions from this analysis[19]. The numbers in the panel after restrictions are applied decline over time and the numbers with missing incomes increase. HBS households aged 60+ in 2015 are represented in earlier years and not in 2016 and 2017, while the restriction to first and second incomes means that younger people in households are not coming into the panel unless they are the main household earners in later years.

[19] Sampling errors (1.96 X $\sigma/\sqrt{n}$) for earnings/incomes are +/- €894/€875 in 2005 and €1,552/€1,502 in 2017.

Table 6.1 Summary statistics for male earnings, 2005-2017

Year	Mean	Standard deviation	N	N missing	Lower Quartile	Median	Upper Quartile	95th Pctile
	€/%	€/%	Obs	Obs	€/%	€/%	€/%	€/%
2005	€37,934	€24,045	2,770	747	€21,999	€33,736	€48,488	€84,415
2006	€39,941	€26,287	2,849	668	€22,683	€35,157	€51,257	€88,665
2007	€41,931	€26,719	2,840	677	€24,248	€36,664	€54,251	€93,397
2008	€43,106	€27,826	2,729	788	€24,553	€37,829	€55,405	€97,412
2009	€39,695	€26,185	2,565	952	€20,679	€36,053	€53,374	€88,186
2010	€38,483	€26,302	2,454	1,063	€19,455	€34,897	€52,039	€87,016
2011	€38,414	€26,406	2,351	1,166	€19,227	€34,524	€52,134	€86,615
2012	€39,504	€28,821	2,307	1,210	€17,287	€34,917	€54,773	€95,653
2013	€39,291	€28,321	2,218	1,299	€17,219	€34,970	€56,152	€92,909
2014	€43,813	€30,274	2,149	1,367	€22,170	€38,140	€59,443	€102,435
2015	€46,692	€33,843	2,125	1,392	€23,311	€39,002	€61,881	€108,641
2016	€46,096	€34,202	2,154	1,363	€22,605	€38,899	€62,277	€109,938
2017	€48,848	€36,149	2,084	1,433	€24,807	€40,466	€65,191	€113,226
2005-17	28.8%	50.3%			12.8%	19.9%	34.4%	34.1%

Table 6.2 Summary statistics for male incomes, 2005-2017

Year	Mean	Standard deviation	N	N missing	Lower Quartile	Median	Upper Quartile	95th Pctile
	€/%	€/%	Obs	Obs	€/%	€/%	€/%	€/%
2005	€37,100	€24,019	2,894	623	€20,724	€32,887	€47,983	€82,939
2006	€39,092	€26,219	2,970	547	€21,532	€34,359	€50,429	€88,022
2007	€41,270	€26,565	2,946	571	€23,414	€35,836	€53,147	€92,068
2008	€42,081	€27,683	2,865	652	€22,964	€36,862	€54,468	€95,780
2009	€38,278	€25,871	2,772	745	€18,358	€34,385	€51,697	€86,099
2010	€36,740	€25,880	2,697	820	€15,983	€32,621	€50,000	€85,543
2011	€36,458	€26,071	2,599	918	€15,404	€32,314	€49,752	€84,506
2012	€37,527	€28,317	2,544	973	€13,866	€32,424	€52,864	€92,828
2013	€37,260	€27,971	2,437	1,080	€13,477	€32,759	€53,185	€89,516
2014	€42,351	€29,901	2,301	1,215	€20,250	€36,204	€57,710	€100,281
2015	€45,029	€33,501	2,271	1,246	€21,219	€37,141	€60,198	€106,888
2016	€45,933	€33,465	2,223	1,294	€21,965	€38,058	€61,355	€109,022
2017	€47,745	€35,784	2,181	1,336	€23,086	€39,191	€63,653	€111,499
2005-17	28.7%	49.0%			11.4%	19.2%	32.7%	34.4%

Figure 6.1 shows male log earnings variance in each cross-section over the 2005-2017 period. The variance of male log earnings doubles during the crisis, which has not been found in any Irish analysis of inequality so far, including research with a gender dimension. The variance

rises steadily from 2008 to peak levels in 2012 which requires further analysis on two questions: What drives this significant increase in male earnings risk during the period? Also, how do welfare and spouse incomes moderate this increase in male earnings risk?

Figure 6.1 Male log earnings variance, 2005-2017

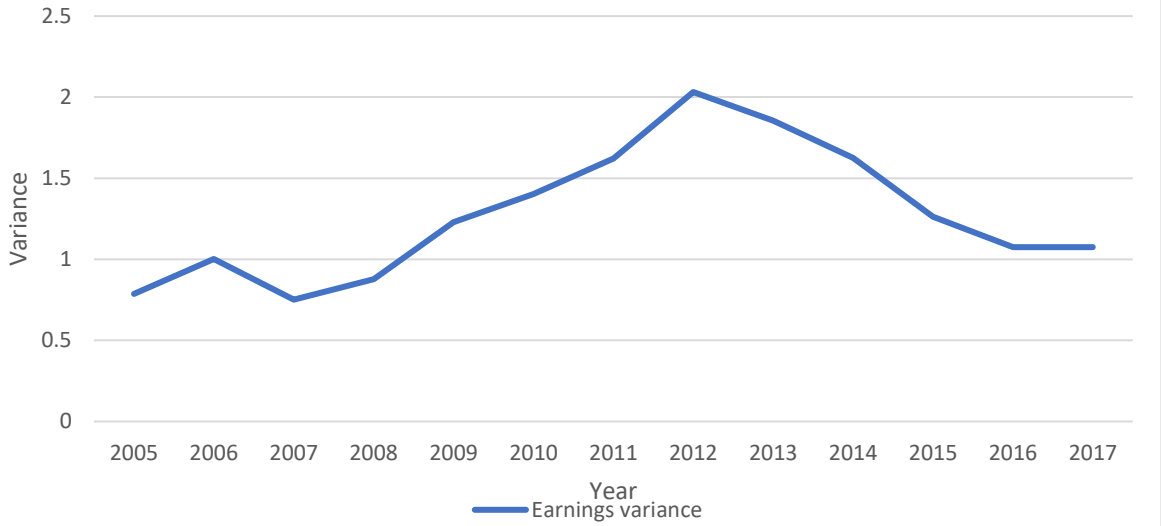

Most of the Irish research to date on incomes and earnings during the crisis has been based on annual SILC cross section data. Researchers focus on inequality of incomes and earnings over time, which has been quite stable as discussed earlier, rather than volatility since it does not require repeated observation of the same people in a panel. One such analysis by Holton and O'Neill (2017) was based on hourly earnings for employees in repeated cross-sections of the SILC survey from 2004 to 2013. My analysis so far has been based on annual earnings of males regardless of their employment status within a defined age band from an administrative income panel. Holton and O'Neill found that low wage employees were reasonably well protected during the crisis which appears to be at odds with the results above. This is likely to be because job losses and reduced hours are reflected at individual level in my income panel. Looking at Table 6.1, reduced hours and pay are reflected in the 20 per cent drop in lower quartile earnings levels from 2008-2010, while earnings for males coming back into the labour market in 2012-13 are also relatively low, resulting in a further fall in earnings at the lower quartile level. Men were particularly affected by job losses due to the collapse of construction.

Figure 6.2 plots the trend in median and lower quartile earnings in my data. The trend in median earnings compares well to Holton and O'Neill's analysis, with a 'hump shaped' pattern of earnings during the crisis. Earnings at this level have recovered since the crisis period. The

trend in lower quartile earnings shown in Figure 6.2 differs from Holton and O'Neill's analysis, which showed stability in the bottom quartile of hourly earnings. In my analysis bottom quartile earnings fell more sharply than median earnings through most of the crisis years. There was a significant increase in male part time and casual employment during the crisis[20], which particularly reduces bottom quartile earnings levels and is not possible to analyse in my data since hours are not available. There are much larger numbers of very small annual earnings amounts in the crisis years, and there is no evidence of non-compliance with minimum wage policy, so part time and irregular employments must explain the differences in earnings in both series.

Figure 6.2 Median and lower quartile values for earnings, 2005-2017

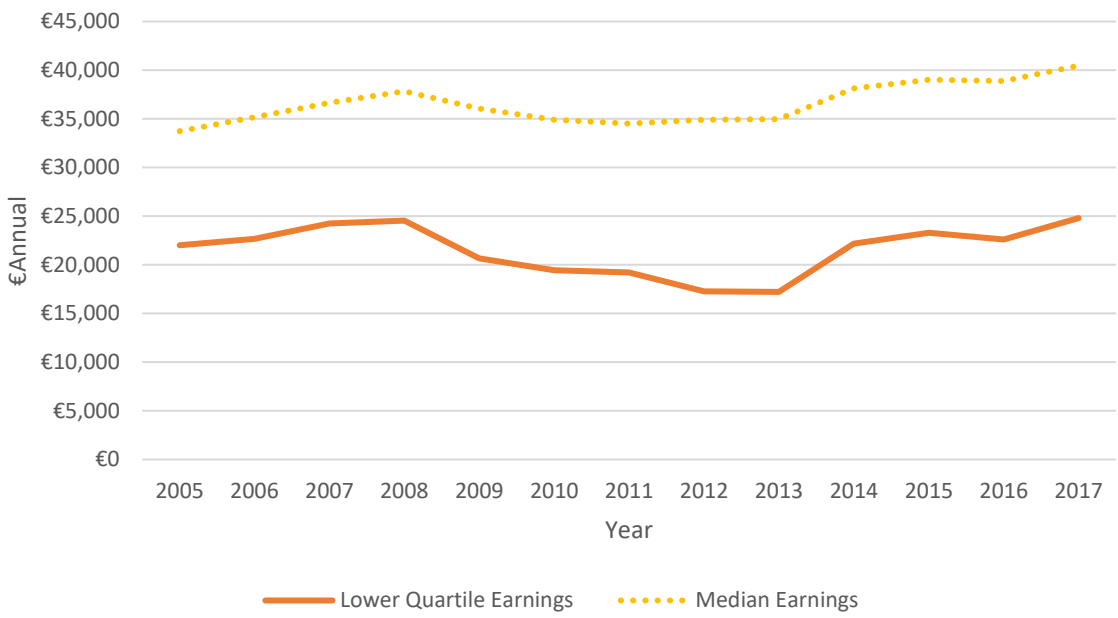

6.4 Methodology

None of the above analysis makes use of the panel aspect of these data. However, exploiting this feature allows us to move beyond standard measures of inequality and focus on transitory inequality and volatility.

Moffitt and Gottschalk (2002) used a simple graphical presentation to illustrate the trend in transitory earnings in the US from 1967-1996, in addition to a more sophisticated econometric model which showed the same trends. The graphical presentation was based on comparisons

[20] https://www.cso.ie/en/releasesandpublications/er/qnhs-es/qnhsemploymentseriesq12016/

of the variance of earnings in a year with covariances in subsequent years. The standard permanent-transitory model is;

$$y_{ia} = \mu_i + \upsilon_{ia} \quad (1)$$

Where y_{ia} is log earnings for individual i at age a, μ_i is a time invariant individual component with variance σ_μ^2 and $\nu_{ia}$is a transitory component with variance σ_ν^2. Assuming the two components are uncorrelated, the cross-sectional variance of log earnings is the sum of the two;

$$V = Var(y_{ia}) = \sigma_\mu^2 + \sigma_\nu^2 \quad (2)$$

Moffitt and Gottschalk's simple presentation was based on the permanent variance of earnings being equal to the covariance of log earnings between a pair of ages sufficiently far apart so that the transitory errors are uncorrelated, so time enters as a parameter of the model;

$$C = Cov(y_{ia,t}, y_{ia,t+n}) = \sigma_\mu^2 \qquad if\ Cov(\nu_{ia,t}, \nu_{ia,t+n}) = 0 \quad (3)$$

An estimate of the transitory variance can be computed from sample data as T = V – C.

However, consider a model in which a factor loading α_t is included with the permanent component:

$$y_{iat} = \alpha_t \mu_{iat} + \upsilon_{iat} \quad (4)$$

The expressions for the variance and covariance in (4) are;

$$Var(y_{iat}) = \alpha_t^2 Var(\mu_{iat}) + Var(\nu_{iat}) \quad (5)$$

$$Cov(y_{iat}, y_{i,a-1,t-1}) = \alpha_t \alpha_{t-1} Var(\mu_{iat}) \quad (6)$$

$$\text{V - C} = Var(\nu_{iat}) + \alpha_t\ (\alpha_t - \alpha_{t-1}) Var(\mu_{iat}) \quad (7)$$

This model includes both transitory and permanent components, and only results in the transitory variance when $\alpha_t = \alpha_{t-1}$.

Shin and Solon (2011) raised issues with both the simple and more formal approaches. Their proposal is to use an earnings volatility measure based on the dispersion in year-to-year earnings changes reflecting both permanent and transitory shocks, which in their view is much more relevant to volatility measurement than variance components models. Based on the

formal model in Equation (4), the variance of the change in log earnings between years *t-2* and *t* is;

$$Var(y_{iat} - y_{ia,t-2}) = (\alpha_t - \alpha_{t-2})^2 Var(\mu_{iat}) + Var(v_{iat}) + Var(v_{ia,t-2}) \quad (8)$$

This dispersion measure is higher when the transitory variance is higher in years t and t-2, so income volatility arising from transitory shocks is fully reflected in the measure. Changes in α_t distort the measure through their impact on the permanent variance, but Shin and Solon find that selection of α closer in time reduces this distortion. The fractional change in the α 's is squared which further reduces the distortion.

Shin and Solon include a random walk income process in the model in Equation (4), which extends it to include the impact of permanent income shocks;

$$y_{iat} = \alpha_t(p_i + \mu_{iat}) + v_{iat} \quad (9)$$

Where μ_{iat} follows a martingale process affecting all subsequent years of income;

$$\mu_{iat} = \mu_{iat-1} + \omega_{iat} \quad (10)$$

The variance measure in Equation (8) is modified to;

$$Var(y_{iat} - y_{ia,t-2}) = (\alpha_t - \alpha_{t-2})^2 Var(p_i + \mu_{iat-2}) + Var(v_{iat}) + Var(v_{iat-2}) + \alpha_t^2 [Var(\omega_{iat}) + Var(\omega_{iat-2})] \quad (11)$$

where $p_i + \mu_{iat-2}$ is described as worker *i*'s permanent human capital in time *t-2*. The main change relative to Equation (8) is the addition of the last term $\alpha_t^2 [Var(\omega_{iat}) + Var(\omega_{iat-1})]$, which is the component of the variance in earnings change that comes from permanent shocks. They argue that permanent shocks are key components of income volatility and are unpredictable. Their volatility measure includes both types of shocks, which is appropriate in this case. Shin and Solon use residual log earnings from a regression on age in these calculations to focus only on the time shifts in volatility. They use the square root of the expression in (8) as a measure of volatility:

$$S = \sqrt{Var[logE_{it} - logE_{it-2}]}$$

Where $logE_{it}$ are the residuals from the log earnings regression.

6.5 Results

I use a first stage Ordinary Least Squares regression to identify and remove the relationship between earnings and age to focus only on time shifts in volatility, in line with Shin and Solon's approach. Age effects could reasonably be thought to be predictable and thus should be removed from any measure of volatility (Jenkins, 2011). I use log earnings as the dependent variable and age and age squared as the independent variables. Figure 6.3 shows log earnings volatility based on the Shin and Solon measure with a one-year lag for the main panel (aged 25-60) and for a further restricted panel aged 30-54. Both Baker and Solon (2003) and Shin and Solon similarly show lower volatility for middle age groups, which reflects higher volatility for younger and older workers. Volatility can be seen to increase rapidly at the start of the crisis and again in 2013-14 as earnings bounce back from their lowest levels. However, volatility falls below pre-crisis levels by the end of the period. Shin and Solon and Moffitt and Gottschalk found volatility to be strongly countercyclical, with particularly large increases in the US recessions of the mid 1970's and early 1980's. All of the US studies show volatility at higher long term levels than before the crisis periods, which is not seen in the data here. The restricted panel shows lower volatility than the full panel.

Figure 6.3 Residual earnings volatility for males aged 25-60 and 30-54, 2006-2017

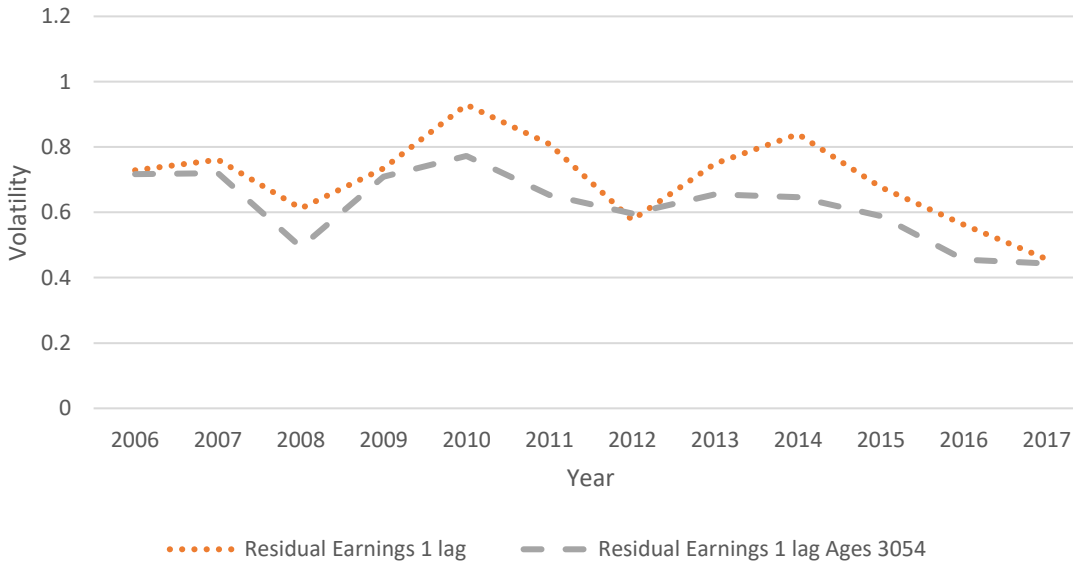

Guvenen et al. (2015) found that the distribution of US volatility or earnings 'shocks' was 'left sided' and not countercyclical on average, i.e. large upward movements in earnings levels are less likely during recessions and large drops are more likely. This finding is based on a similar

administrative dataset for US annual earnings to my panel, which has a much larger left tail of small/irregular annual earnings than surveys such as the PSID (Moffitt, 2021). Figure 6.4 presents the distribution of year-to-year earnings changes in the panel, which can also be interpreted as shocks. Figure 6.4 shows very strong falls in earnings at the 10th percentile of changes in log earnings during 2008-2012, and falls at the 25th percentile are also quite pronounced. Pay increases at higher levels were supressed during this period, until a 'rebound' in 2014, all of which is consistent with Guvenen et al.'s findings on skewness. The increases at the 90th percentile in 2014 are concentrated among formerly low earners. To illustrate this movement, the bottom decile level of earnings was above €10,000 until 2009 and resumed above this level from 2014; it fell to its lowest points at 2012 and 2013 at €6,343 and €6,754 respectively. This change from 2012/13 to 2014 reflects combined improvements in hours and earnings at the lowest levels of 2013 earnings. The spread of the distribution of changes from 2015 on is more symmetric, suggesting more stability in the labour market since the recovery from the crisis.

Figure 6.4 Quantiles of age-adjusted change in male log earnings, 2006-2017

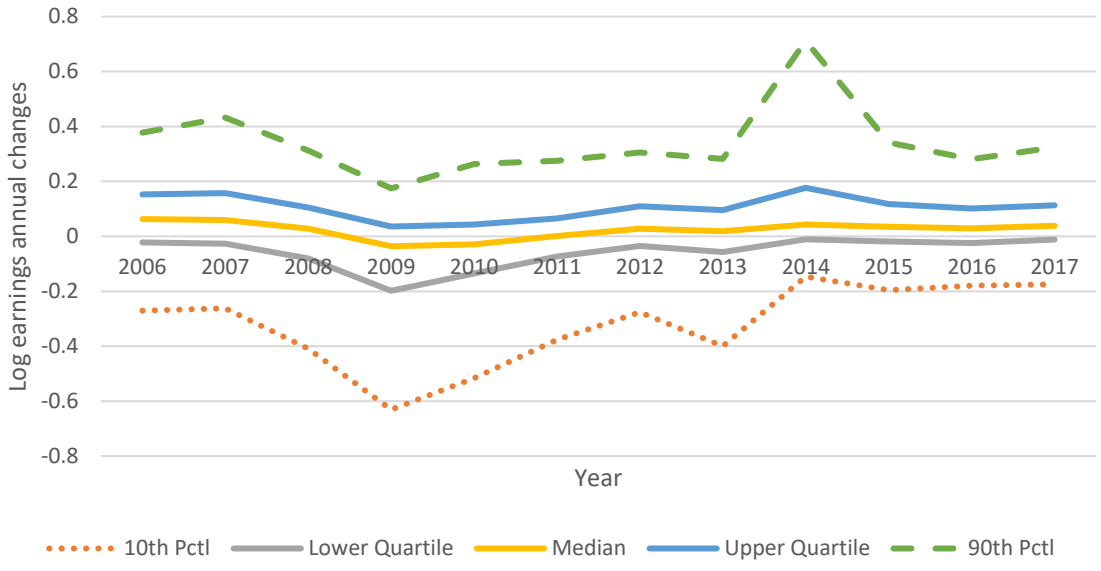

The Shin and Solon earnings volatility analysis is presented separately for males with low (below Leaving Certificate) and high (third level) education in Figure 6.5 below. The low education group experienced higher levels of volatility than average throughout the period, particularly during the crisis, which is at odds with Holton and O'Neill's findings of relative stability of earnings for low education/low paid workers during the crisis. My earnings measure picks up volatility due to lost hours and jobs in addition to the volatility in hourly pay in their

analysis. Volatility for the high education group also shows countercyclical movement, but to a lesser extent than for the low education group. Also, income volatility for the high education group is lower than for low education workers for the period after the crisis.

Figure 6.5 Earnings volatility for High and Low Education groups, 2006-2017

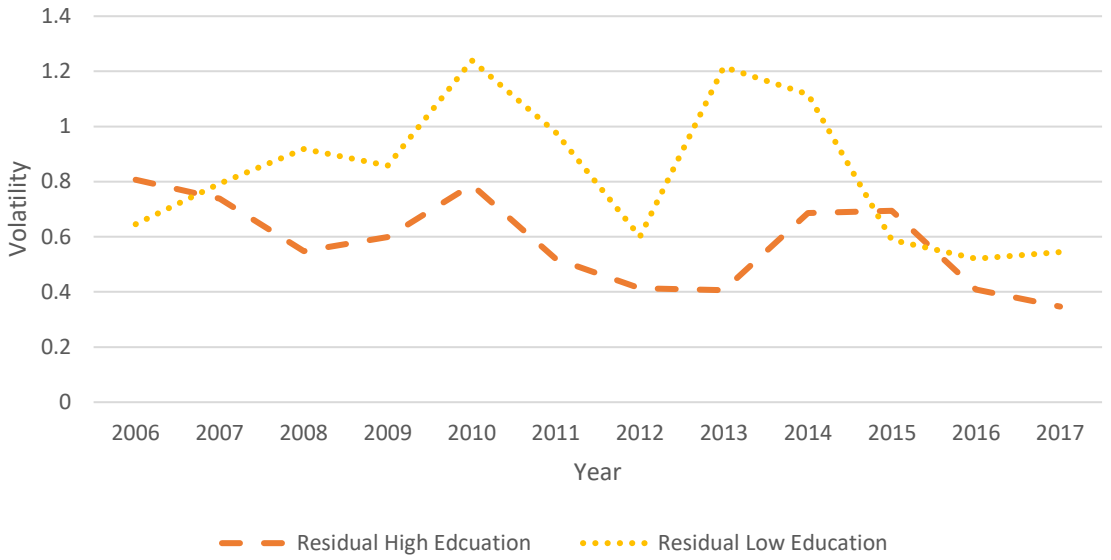

A drawback of the Shin and Solon analysis is that zeros cannot be included due to the log transformation used. This makes it impossible to examine the contribution of different income components, when one of these are zero. Such a comparison is needed to examine the extent of insurance provided by either the welfare system or spousal labour supply. Shin and Solon however, also presented an alternative approach to the log transformation which allows inclusion of zeros in the analysis where panellists have no income from a particular source for the year in question.

To carry out an analysis of the contribution of each income source to volatility, the same sample size is required across all three income measures (male earnings, male incomes and household incomes) each year. The resulting sample when zeros are included instead of missing values for male earnings and incomes starts at 3,100 in 2005 and falls to 2,300 in 2017[21]. Shin and Solon also reintroduced outliers as part of this approach to give a complete decomposition of volatility, and year to year differences in income levels were used instead of logs.

[21] Results from a balanced panel of respondents aged 25-60 in 2015 are included in Appendix 6. The resulting sample is much lower but a similar trend in volatility and insurance to that found in the unbalanced panel is also observed here.

To examine the role of insurance against shocks I estimate the volatility measure for successively broader measures of income. I start with male earnings which now includes zeros for those with no earnings. I then estimate the volatility measure for male income by adding in welfare payments accruing to men. Finally, I estimate household volatility by adding in spousal income.

Figure 6.6 is presented using this approach and shows a pattern of 'insurance' for male welfare from 2008 onwards. Volatility falls as we look at each broader definition of income, consistent with the additional income source offering insurance against risk. Throughout the period volatility based on household income is approximately half that based on male earnings only. This may be due to substitution of spouse's labour supply, since there is an increased reliance on spouse incomes in households and Figure 6.6 also shows a flatter trend in volatility for household incomes than earnings during the crisis. Dynan el al (2012) found limited evidence of substitution in the US based on a similar analysis.

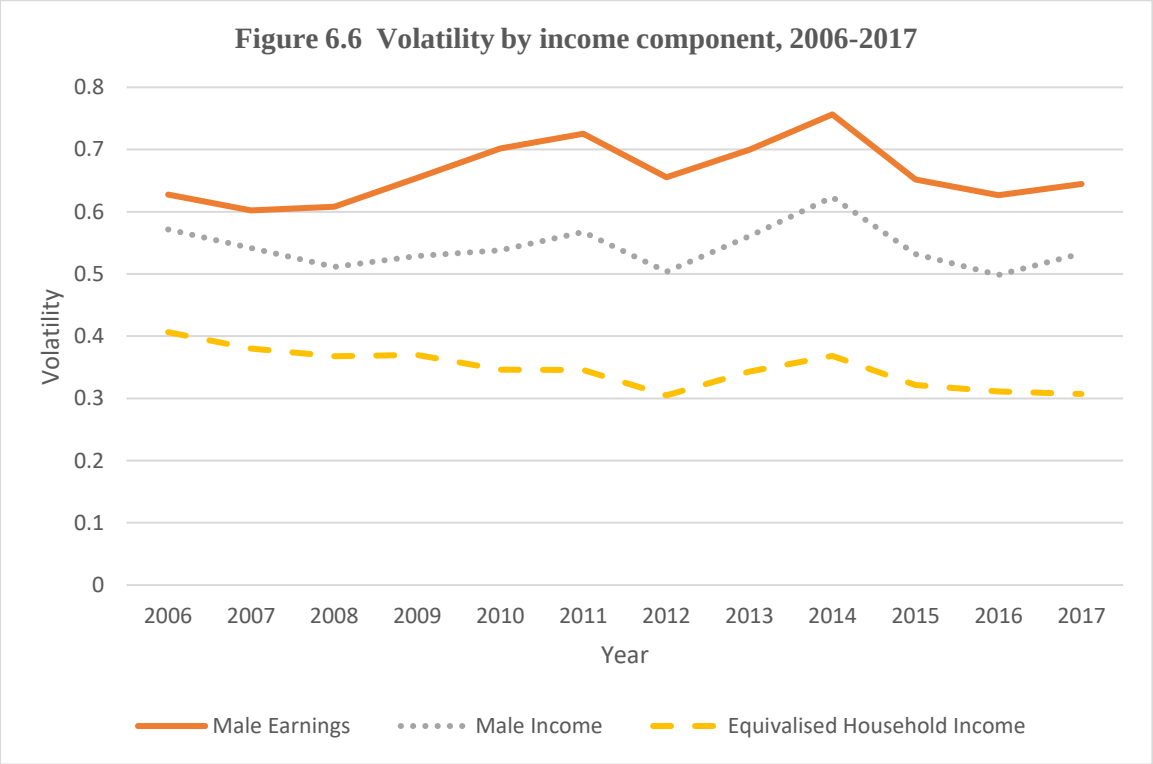

Figure 6.7 presents the same analysis as Figure 6.6 but excluding the top 1 per cent of incomes only, which shows a very similar pattern of volatility to Figure 6.6 and thus limited influence of top incomes in volatility results. Figure 6.8 additionally excludes the zero values for male earnings and incomes from the initial analysis in Figure 6.6 and is included to demonstrate the difference in results from the alternative Shin and Solon approach. When zeros are excluded,

annual samples are different for each of the income measures. Gaps in volatility between the three income measures are much lower on this basis, since full income risks are not measured for male earnings and income, which makes the insurance comparison of household incomes, male income and earnings meaningless.

Figure 6.7 Volatility by income component excluding top incomes, 2006-2017

Figure 6.8 Volatility by income component excluding top incomes and zeros, 2006-2017

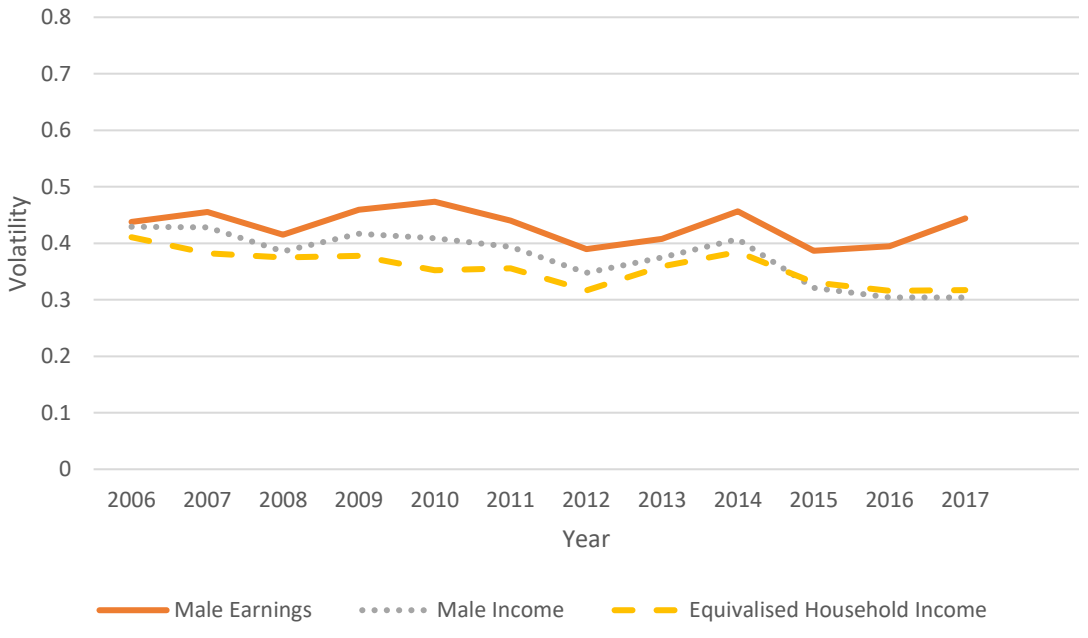

6.6 Conclusions and Discussion

The existing literature on Irish inequality has found that it has been quite stable over the long term, including during the "Great Recession". In this paper I use a new linked panel data set to examine risk and volatility over this period and the role of the welfare system and spousal labour supply as insurance against unanticipated shocks.

Analysis of volatility based on administrative panel data reveals that male earnings risks increased significantly during the "Great Recession", which has been hidden to date by the use of cross-section survey data. A combination of job losses, pay cuts and hours reductions are reflected in changes in annual earnings data, which combine to drive a significant increase in volatility in 2008-2010. The increase in volatility is mainly in low income and low education groups, who had far more volatile incomes than median or high earners. A second peak in earnings volatility is identified in 2013-2014 as employment, hours and earnings simultaneously improved rapidly for low earner males, which is unusual in an international context. Many males who lost their jobs during the "Great Recession" re-entered the labour market via casual and part time employment in construction and other recovering sectors, which shows up as rebound volatility in the data. Earnings volatility falls back to below pre-crisis levels after 2014, whereas many US researchers found a 'level shift' upwards in volatility and related measures after crises. Male earnings volatility remains relatively high in 2015-17 based on international comparisons despite the fall from peak levels. The availability of social welfare incomes and household linkages from HBS facilitates the analysis of the 'insurance' provided by male welfare payments and spouses' incomes. Both combine to result in flatter, lower levels of volatility. Throughout the sample household volatility is half that of male earnings. The partial insurance offered by welfare and spousal labour supply helps mitigate the impact of unanticipated shocks allowing households to partially smooth consumption in the face of these shocks.

Appendix 6 Results from a balanced panel aged 25-60 in 2015

Table 6.3 Summary statistics for male earnings in balanced panel, 2005-2017

Year	Mean	Standard deviation	N	N missing	Lower Quartile	Median	Upper Quartile	95th Pctile
	€/%	€/%	Obs	Obs	€/%	€/%	€/%	€/%
2005	€37,734	€23,300	1,049	261	€23,192	€33,984	€48,254	€79,881
2006	€40,405	€26,372	1,080	230	€23,957	€36,363	€52,141	€89,011
2007	€42,174	€26,011	1,107	203	€24,905	€37,625	€55,283	€89,467
2008	€43,715	€27,187	1,093	217	€25,694	€38,808	€57,058	€94,812
2009	€40,431	€26,069	1,046	264	€21,446	€36,893	€55,026	€87,577
2010	€40,094	€26,307	1,014	296	€21,663	€36,787	€54,732	€87,377
2011	€39,880	€27,119	1,013	297	€20,155	€35,718	€53,730	€90,700
2012	€41,239	€30,005	1,017	293	€17,932	€36,016	€57,489	€101,590
2013	€40,857	€29,112	1,014	296	€18,135	€35,693	€58,679	€95,985
2014	€44,346	€31,007	1,013	297	€21,830	€38,040	€61,363	€101,249
2015	€47,700	€35,061	1,026	284	€23,401	€39,309	€63,210	€112,762
2016	€48,039	€35,555	1,050	260	€23,120	€39,474	€64,765	€112,649
2017	€49,951	€37,475	1,064	246	€24,901	€40,911	€66,402	€118,247
2005-17	32.4%	60.8%			7.4%	20.4%	37.6%	48.0%

.

Table 6.4 Summary statistics for male incomes in balanced panel, 2005-2017

Year	Mean	Standard deviation	N	N missing	Lower Quartile	Median	Upper Quartile	95th Pctile
	€/%	€/%	Obs	Obs	€/%	€/%	€/%	€/%
2005	€36,939	€23,288	1,093	217	€21,875	€33,176	€47,563	€79,563
2006	€39,493	€26,358	1,124	186	€22,237	€35,291	€51,505	€86,048
2007	€41,546	€25,900	1,145	165	€24,479	€36,900	€54,841	€88,436
2008	€42,780	€27,129	1,141	169	€24,670	€37,764	€55,993	€93,963
2009	€38,816	€25,846	1,133	177	€18,866	€35,058	€53,000	€84,845
2010	€38,067	€26,072	1,117	193	€17,613	€34,154	€52,834	€85,000
2011	€37,781	€26,868	1,118	192	€17,050	€33,089	€52,075	€86,615
2012	€39,285	€29,507	1,115	195	€14,613	€33,641	€55,175	€98,085
2013	€38,800	€28,841	1,111	199	€13,987	€33,239	€55,589	€93,400
2014	€42,887	€30,600	1,086	224	€19,760	€35,891	€59,099	€100,281
2015	€45,824	€34,659	1,104	206	€21,006	€37,127	€61,491	€110,075
2016	€46,933	€35,020	1,108	202	€21,559	€38,029	€63,453	€112,257
2017	€48,586	€37,088	1,123	187	€22,620	€39,303	€65,000	€116,305
2005-17	31.5%	59.3%			3.4%	18.5%	36.7%	46.2%

Figure 6.9 Volatility by income component from balanced panel, 2006-2017

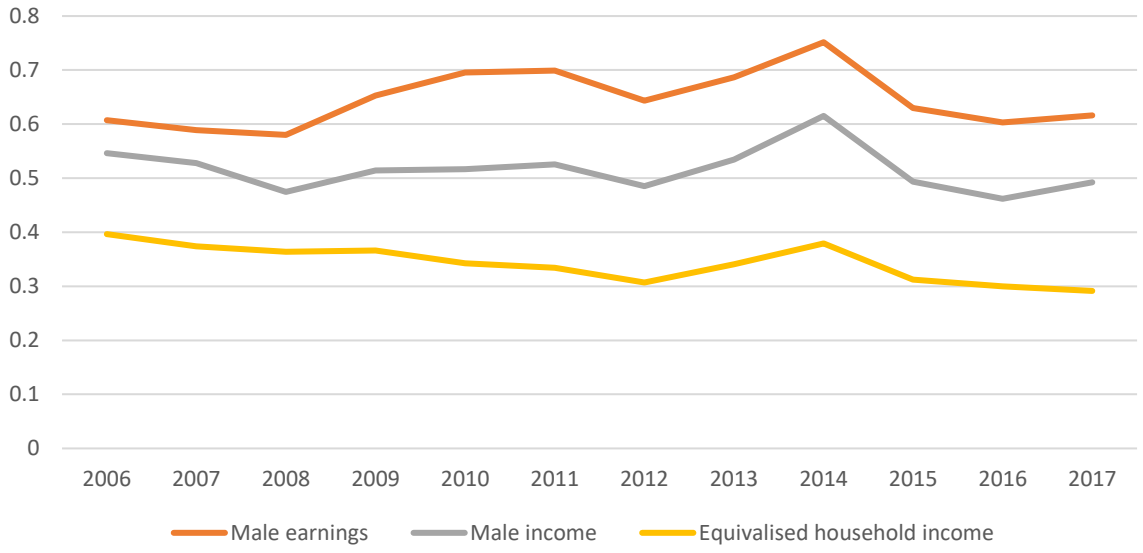

7 Future directions for research

There are some formal models in the literature which could address the validity of the PIH in the Irish context, though they may require a longer income panel than I have for Chapter 6. The approach of Aaberge and Mogstad (2015) in measuring 'life-cycle bias' looks particularly relevant in Ireland, since lifetime incomes may not have been increasing as quickly as current incomes. The approaches of Scholz et al. (2006) and Crawford and O'Dea (2014) in measuring oversavings by pensioners based on lifetime earnings could also offer insights into outcomes for Irish pensioners.

The consumption data in the HBS could be used to extend the Chapter 6 analysis in novel ways. Scarring in 2015 consumption for people who suffered a previous income shock has been mentioned as an addition to the current literature, since availability of consumption with an income panel is unusual. There is enough information in the panel to also look at a group of people with low 'replacement rates' from pensions and how their consumption compares to other retirees. This analysis is likely to have some of the identification issues discussed in Chapter 4, since dates of retirement are not available and proxies (e.g., receipt of State pension) have to be used. Also, Browning and Crossley (2001) outlined some useful formal approaches to modelling consumption that account for durability and investment decisions. Housing is such an important aspect of Irish consumption that it could be looked at in different ways to the rental equivalence approach in alternative consumption models. Owner occupied housing could also be treated as an asset, investment or bequest in a consumption model, and the bequest approach could possibly explain the apparent low level of savings in the HBS.

Finally, the analysis of consumption inequality and poverty in Chapters 3 and 4 could be expanded in line with the vast literature on incomes for Ireland when results from the upcoming annual HBS become available. Increasing measurement error in HBS equivalent surveys for expenditure (Blundell and Etheridge, 2010; Aguiar and Bils, 2015) has limited these approaches in other countries, but the Irish HBS data appears to be more robust.

References

Aaberge, R. and Mogstad, M., (2015). 'Inequality in current and lifetime income', Social Choice and Welfare, Vol. 44(2), pp 217–230

Aguiar, M. and Bils, M., (2015), 'Has Consumption Inequality Mirrored Income Inequality?', American Economic Review, Vol. 105(9) pp 2725-2756

Aguiar and Hurst, E. (2005), 'Consumption vs expenditure', Journal of Political Economy, Vol. 113 (5), pp 919-948

Aguiar and Hurst (2008), 'Deconstructing Lifecycle expenditure', NBER Working Paper 13893

Aguila, E., Attanasio, O., and Meghir, C. (2011), 'Changes in consumption at retirement', The Review of Economics and Statistics, Vol. 93(3), pp 1094-1099

Attanasio Berloffa, G., Blundell, R. and Preston, I. (2002), 'From Earnings Inequality to Consumption Inequality', Economic Journal Vol 112 No. 478, pp 52-59

Attanasio, Banks, J., Meghir and Weber, G. (1999), 'Humps and bumps in lifetime consumption', Journal of Business and Economic Statistics, Vol.17 (1), pp 22-35

Attanasio and Weber (1995), 'Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey', Journal of Political Economy, Vol.103 (6), pp 1121-1157

Baker, M. and Solon, G. (2003), "Earnings Dynamics and Inequality among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records." Journal of Labor Economics, 21(2): 267-288.

Banks, Blundell and Tanner, S. (1998), 'Is There a Retirement-Savings Puzzle?', American Economic Review , Vol. 88, No. 4. (Sep., 1998), pp 769-788

Banks, Emerson, C. And Oldfield, Z. (2004), 'Not so brief lives: Longevity expecations and wellbeing in retirement', in Iain Stewart and Romesh Vaitlingham, eds., *Seven ages of man*

and woman: A look at life in Britain in the second Elizabethan era, Economic and Social Research Council

Battisin, E., Brugiavini, A., Rettore, E. and Weber (2009), 'The retirement consumption puzzle; evidence from a regression discontinuity approach', American Economic Review, Vol. 99 (5), pp 2209-2226

Bernheim, B.D., Skinner, J. and Weinburg, S. (2001), 'What accounts for the variation in retirement wealth among U.S. households?', American Economic Review, Vol. 91 (4), pp 832-857

Bissonnette, L., Hurd, M. and Michaud P.-C. (2017). 'Individual survival curves comparing subjective and observed mortality risks'. Health Economics, pp 1–19.

Blau, D. (2008), 'Retirement and consumption in a life-cycle model', Journal of Labor Economics, Vol 26 (1), pp 35-71

Blundell, Pistaferri, L., and Preston, 'Consumption inequality and partial insurance', American Economic Review, Vol. 98 (5), pp 1887-1921

Blundell and Preston (1998), 'Consumption Inequality and Income Uncertainty', Quarterly Journal of Economics, Vol. CXIII (2) pp 603-641

Blundell and Etheridge, B. (2010) 'Consumption, income and earnings inequality in Britain', Review of Income Dynamics, 2010 pp 1094-2025

Blundell, Pistaferri and Saporta-Eksten, I. (2016) "Consumption Inequality and Family Labor Supply", American Economic Review, 106(2): pp 387-435

Brewer, M. and O'Dea, C. (2012), 'Measuring living standards with income and consumption; evidence from the UK', IFS Working Paper W12/12 March 2012

Brewer, Etheridge and O'Dea (2017), 'Why are households that report the lowest incomes so well off?', Economic Journal, Vol. 127; pp 24-49

Browning, M., and Crossley, T.F., (2001) 'The Life-Cycle Model of Consumption and Saving', Journal of Economic Perspectives, Vol. 15 (3), pp 3-22

Caballero, R. (1990), 'Consumption Puzzles and Precautionary Savings', Journal of Monetary Economics, Vol. 25, pp 113-136

Callan, T., Bercholz M. and Walsh, J.R., (2018), 'Income growth and income distribution: a long run view of Irish experience', Budget Perspectives 2019 conference paper, ESRI.

Callan, Doorley, K. and Savage, M. (2018), 'Inequality in EU Crisis Countries: How Effective Were Automatic Stabilisers?', IZA Discussion Paper No. 11439, March 2018.

Campbell, J. and Mankiw, N.G. (1989), 'Consumption, income and interest rates; reinterpreting the time series evidence', in Olivier J. Blanchard and Stanley Fischer (eds.), NBER macroeconomics annual 1989, Cambridge Mass. MIT Press, pp 185-216

Campbell and Mankiw (1991), 'The response of consumption to income; a cross-country investigation', European Economic Review, Vol. 35, pp 715-721

Carroll, C. and Summers, L. (1991), 'Consumption growth parallels income growth; some new evidence', in B. Douglas Bernheim and John B. Shoven (eds.),'National Saving and Economic Performance', Chicago, Chicago university Press for NBER, pp. 305-343

Carroll (1992), **'**The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence', Brookings Papers on Economic Activity, Vol. 1992, No. 2 (1992), pp. 61-156

Carroll and Toche, P. (2019), 'A Tractable Model of Buffer Stock Saving', Johns Hopkins University, April 10 2019

Collins, M., 'Modelling the Distributive Impact of Indirect Taxation Changes in the Republic of Ireland', Nevin Economic Research Institute working paper June 2014

Crawford, R. and O'Dea (2014), 'Cash and Pensions: Have the elderly in England saved optimally for retirement?', IFS Working Paper W14/22

Cutler, D. and Katz, L. (1992), 'Rising Inequality? Changes in the distribution of Income and Consumption in the 1980's', American Economic Review, Vol. 82 (2) pp 546-551

Deaton, A. (1985), 'Panel data from a time-series of cross sections', Journal of Econometrics, Vol. 30, pp 109-126

Deaton (1991), 'Savings and Liquidity Constraints', Econometrica LIX (1991), pp 1221-1248

Deaton and Paxson, C. (1994), 'Intertemporal Choice and Inequality', Journal of Political Economy, Vol. 102 No. 3 (June 1994), pp 437-467

Dooley, J. and Madden, D. (2022), 'Ireland's Post-Crisis Recovery, 2012-2019: Was It Pro-Poor?', The Economic and Social Review, Vol. 53(2, Summer), pp 137-172

Doris, A., O'Neill, D and Sweetman O. (2015). "Wage Flexibility and the Great Recession: the Response of the Irish Labour Market", IZA Journal of European Labor Studies, Vol. 4, p. 18

Dynan, K., Elmendorf, D. and Sichel, D. (2012), 'The Evolution of Household Income Volatility', The B.E, Journal of Economic Analysis and Policy Advances, 12/2

Elder, T. E. (2013), 'The predictive validity of subjective mortality expectations: evidence from the Health and Retirement Study', Demography 50(2), pp 569–589.

Fernandez-Villaverde, J. and Krueger, D. (2007), 'Consumption over the life cycle; facts from the Consumer Expenditure Survey', The Review of Economics and Statistics, Vol. 89 (3), pp 552-565.

Fisher, J., Johnson, D., Marchand, J., Smeeding, T. and Boyle Torrey, B. (2006), 'The retirement consumption conundrum; evidence from a consumption survey', Centre for Retirement Research Working Paper 109

Flavin, M. (1981), 'The adjustment of consumption to changing expectations about future income', Journal of Political Economy, Vol. 89, pp 974-1009

Flavin (1993), 'The excess smoothness of consumption; identification and interpretation', Review of Economic Studies, Vol. 60 (3), pp 651-666

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

Gan, L., Hurd, and McFadden, D. (2005), 'Individual subjective survival curves', in David Wise ed,, *Analyses in the Economics of Aging*, The University of Chicago Press

Gottschalk, P. and Moffitt, R.A. (1994), 'The growth of earnings instability in the U.S. labor market', Brookings Papers on Economic Activity, 2-1994, pp. 217–72.

Guvenen, F., Okzan, S. and Song, J. (2015), 'The Nature of Countercyclical Income Risk', Journal of Political Economy, Vol. 122 No. 3, pp. 621-660

Haider, S.J. (2001), ‘Earnings instability and earnings inequality of males in the United States: 1967–1991’, Journal of Labor Economics, 19, pp. 799–836.

Haider and Stephens, M. (2007), ‘Is there a retirement consumption puzzle? Evidence using subjective retirement expectations’, Review of Economics and statistics, Vol. 66(1), pp 1-7

Hall, R. (1978), ‘Stochastic Implications of the Life Cycle-Permanent Income Hypothesis; Theory and Evidence’, Journal of Political Economy, Vol. 86, No. 6 pp 971-987

Hall (1988), ‘Intertemporal substitution in consumption’, Journal of Political Economy, Vol. 96, pp 339-357

Hall and Mishkin, F. (1982), ‘The sensitivity of consumption to transitory income; estimates from panel data on households’, Econometrica, Vol. 50 pp 461-481

Hamermesh, D. (1984), "Consumption During Retirement: The Missing Link in the Life-Cycle Hypothesis." Review of Economics and Statistics, February, 1984, 66(1), pp 1-7.

Hamermesh (1985) ‘Expectations, life expectancy and economic behaviour’, Quarterly Journal of Economics, Vol. 10(2), pp 389-408

Hayashi, F. (1985), ‘The effect of liquidity constraints on consumption; a cross-sectional analysis’, Quarterly Journal of Economics, Vol. 100, pp 183-206

Heckman, J. (1971), ‘Three essays on the supply of labor and the demand for market goods’, Ph.D. dissertation, Princeton University

Heckman (1974), ‘Life Cycle Consumption and Labour Supply; an explanation of the relationship between income and consumption over the Life Cycle’, American Economic Review, Vol. 64(1), pp 188-194

Holton, N. and O’Neill (2017), ‘The Changing Nature of Irish Wage Inequality from Boom to Bust’, The Economic and Social Review, Vol. 48, No. 1, Spring, 2017, pp. 1-26

Hurd (2009), ‘Subjective probabilities in household surveys’, Annual Review of Economics 1, pp 543-562

Hurd and McGarry, K. (1995). ‘Evaluation of the Subjective Probabilities of Survival in the Health and Retirement Study’. The Journal of Human Resources 30, pp 268-292

Hurd and McGarry (2002), 'The Predictive Validity of Subjective Probabilities of Survival', Economic Journal 112(482), pp 966–985

Hurd, Smith, J. and Zissimopoulos, J. (2004), 'The effects of subjective survival on retirement and social security claiming', Journal of Applied Econometrics 19(6), pp 761– 775

Hurd and Rohwedder, S. (2003), 'The retirement consumption puzzle; anticipated and actual declines in spending at retirement', NBER Working Paper 9586

Hurd and Rohwedder (2008), 'The retirement consumption puzzle; actual spending change in panel data', NBER Working Paper 13939

Hurst (2008), 'The Retirement of a consumption puzzle', NBER Working Paper 13789

Jappelli, T. and Pistaferri (2008), 'Financial integration and consumption smoothing', CSEF Working Paper 200

Krueger and Perri, F. (2006), 'Does Income Inequality Lead to Consumption Inequality? Evidence and Theory', Review of Economic Studies, 2006, v73(1,Jan), pp 163-193.

Madden, D., (2014), 'Winners and Losers on the Roller-Coaster: Ireland, 2003-2011', The Economic and Social Review, Vol.45(3), pp 405-421

Madden, D., Nolan, A. and Nolan, B. (2005), 'GP Reimbursement and Visiting Behaviour in Ireland', Health Economics, Vol. 14, No. 10, pp 1047-1060

Meghir and Pistaferri (2010), 'Earnings, consumption and lifecycle choices', Working Paper 15914, Cambridge MA, National Bureau of Economic Research, http://www.nber.org/papers/w15914

Meyer, B. and Sullivan, J. (2003), 'Measuring the Well Being of the Poor using Income and Consumption', NBER Working Paper No. 9760, June 2003

Meyer and Sullivan (2009), 'Five Decades of Consumption and Income Poverty', NBER Working Paper No. 14827 March 2009

Meyer and Sullivan (2013), 'Consumption and Income Inequality in the US Since the 1960s', unpublished paper dated 2nd April 2013

Miniaci, R. Monfardini, C. and Weber (2003), 'Is there a retirement consumption puzzle in Italy?', Institute for Fiscal Studies Working Paper WP03/14

Modigliani, F. (1966), 'The Life Cycle Hypothesis of Saving, the Demand for Wealth and the Supply of Capital', Social Research, Vol 33 No. 2, pp 160-217

Moffitt (2021), 'Trends in earnings volatility among US men', 'The Reporter' No. 4 December 2021, National Bureau of Economic Research

Moffitt and Gottschalk (2002), 'Trends in the transitory variance of earnings in the United States', Economic Journal, 112, pp. C68–73

Moffitt and Gottschalk (2012), 'Trends in the Transitory Variance of Male Earnings: Methods and Evidence', The Journal of Human Resources, Winter 2012, Vol. 47, No. 1, pp. 204-236

Nagatani, K. (1972), 'Life cycle saving; theory and fact', American Economic Review, Vol. 62 (3), pp 344-353

Nivakoski, S (2020), 'Wealth and the effect of subjective survival probability', Journal of Population Economics, 33, 633-670.

Nolan, Callan, Whelan, C., Whelan, B. and Williams, J., 'Poverty in the 1990s: evidence from the 1994 Living in Ireland Survey', ESRI research series 170/CPA #23 (1996)

O'Dea and Preston (2014), 'Can We Measure Who Loses Most from Public Service Spending Cuts?', Budget Perspectives 2015, Paper 2., ESRI June 2014.

O'Dea, and Sturrock, D. (2018), 'Subjective expectations of survival and economic behaviour', IFS Working Paper W18/14

O'Neill and Sweetman (2001), 'Inequality in Ireland 1987-1994: A Comparison using Measures of Income and Consumption', Journal of Income Distribution, 10: pp 23-39

Parzen, E. (1962) 'On Estimation of a Probability Density Function and Mode', The Annals Of Mathematical Statistics 33, no. 3, pp 1065--1076.

Perozek, M. (2008), 'Using subjective expectations to forecast longevity: Do survey reposndents know something we don't know?, Demography, Vol. 45(1), pp 95-113

Piketty, T. (2014), ‘Capital in the twenty-first century’, The Belknap Press of Harvard University Press, 2014

Poterba, J. (1991), ‘Is the gasoline tax regressive?’, Tax Policy and the Economy, ed. David Bradford, MIT press pp 145-164

Redmond, P. and McGuinness, S. (2022), ‘Consumption in Retirement: Heterogeneous Effects by Household Type and Gender’, Population Ageing Vol. 15, pp 473–491

Regan, M. and Kakoulidou, T. (2022), ‘How important are the unit of analysis and equivalence scales when measuring income poverty and inequality? Evidence from Ireland’, ESRI Working Paper 721

Roantree, B. (2020), ‘Understanding income inequality in Ireland’, 2020 Barrington Prize winner, Statistical and Social Inquiry Society of Ireland Journal

Rosenblatt, M. (1956), ‘Remarks on Some Nonparametric Estimates of a Density Function’, The Annals of Mathematical Statistics, 27, pp 832-837.

Scholz, J, Seshadri, A. and Khitatrakun, S. (2006): “Are Americans Saving ‘Optimally’ for Retirement?”, Journal of Political Economy, 114, pp 607–643

Shin, D. and Solon, G. (2011), ‘Trends in men’s earnings volatility: what does the Panel Study of Income Dynamics show?’, Journal of Public Economics, 95, pp. 973–82

Slesnick, D. (1993), ‘Gaining Ground: Poverty in the Postwar United States’, Journal of Political Economy, Volume 101 No. 1 (Feb. 1993), pp 1-38

Smith, S. (2006), ‘The retirement consumption puzzle and involuntary early retirement; evidence from the British Household Panel Survey, Economic Journal, Vol. 116, pp 139-148

Silverman, B. (1986), ‘Density Estimation for Statistics and Data Analysis’, Monographs on Statistics and Applied Probability, Chapman and Hall, London

Tanner (1998), ‘How much do consumers spend? Comparing the FES and National Accounts’ in Banks, James and Johnson, Paul (eds.), ‘How Reliable is the Family Expenditure Survey? Trends in Incomes and Expenditures Over Time’, IFS Report 57

Thurow, L. (1969), ‘The optimum lifetime distribution of consumption expenditures’, American Economic Review, Vol. 59, pp 324-330

Tobin, J. (1967), ‘Life-cycle saving and balanced growth’ in W. Fellner et al. (eds.), ‘Ten economic studies in the tradition of Irving Fisher’, New York, Wiley, pp 327-368

Wu, S., Stevens, R. and Thorp, S. (2015).’‘Cohort and target age effects on subjective survival probabilities: Implications for models of the retirement phase’, Journal of Economic Dynamics and Control 55, pp 39–56.