

REFERENCE POINTS IN THE WILD: APPLICATIONS USING HOUSEHOLD DATA

A THESIS SUBMITTED TO THE UNIVERSITY OF MANCHESTER
FOR THE DEGREE OF DOCTOR OF PHILOSOPHY
IN THE FACULTY OF ECONOMICS

2020

Student id: 8025560

School of Social Science

Contents

Abstract	8
Declaration	10
Copyright	11
Acknowledgements	12
1 Subjective Financial Situation as Reference Points	21
1.1 Introduction	21
1.2 Literature Review	24
1.3 Data and SFS Variables	26
1.4 Model	37
1.5 Methodology	43
1.6 Results	45
1.6.1 Unfulfilled Expectations	46
1.6.2 Factor Analysis	50
1.7 Conclusion	56
A.1 Proof with Uncertainty	59
A.2 Robustness Checks	63
2 Making Up for Losses: Estimation of Marginal Utility with Reference-Dependent Preferences	67
1 Introduction	67
2 Data	70
3 Theoretical Framework	76
4 Results	78
4.1 Decomposing GHQ	80

5	Conclusion	83
3	Living in Your Parents' Shadow: Parents' Earnings as an Aspiration for Children	84
1	Introduction	84
2	Data and Earnings Derivation	88
3	Intergenerational Mobility in Contemporary England	96
3.1	Elasticities Across the Distribution	100
4	Unobserved Aspirations Model	106
4.1	Testing the Model	110
5	Conclusion	114
A.1	Robustness Checks for Model Tests	116

Word Count: 41,879

List of Tables

1.1	Economic and demographic variables in the BHPS	27
1.2	Current Financial Situation Summary Statistics	30
1.3	Expected Financial Situation Summary Statistics	30
1.4	Change in Financial Situation Summary Statistics	31
1.5	Change in Financial Situation Summary Statistics	31
1.6	Reason Financial Situation Changed	33
1.7	Effect of Reference Points on Savings	47
1.8	SFS interactions summary statistics	48
1.9	Effect of Unfulfilled Expectations on Savings	49
1.10	Factor Analysis of SFS	51
1.11	Factor Correlations with Original Variables	51
1.12	Cross-Correlations of SFS Variables	52
1.13	PCA Analysis	52
1.14	Disagreements in SFS Within Households	54
1.15	PCA Analysis	55
A1	Effect of Reference Points on Savings	65
A2	Effect of Reference Points on Savings	66
2.1	Examples of Question Types from GHQ Questionnaire	72
2.2	Food Expenditure Bands for Households in BHPS/USoc	74
2.3	Economic and demographic variables in the combined BHPS/USoc dataset	75
2.4	Effect of Reference Points on Marginal Utility of Consumption	79
2.5	Bayliss et al. (2014) Affect Decomposition	81
2.6	Effect of Reference Points on Marginal Utility of Consumption	82
3.1	Demographic variables in the LSYPE (Children)	89
3.2	Demographic variables in the LSYPE (Parents)	90
3.3	Life-Cycle Profile of Earnings (BHPS/USoc Data)	94

3.4	Transition Matrix	96
3.5	Earnings Persistence	98
3.6	Earnings Persistence with Covariates	99
3.7	IEEs across the distribution (unjoined splines, without covariates)	102
3.8	Earnings Persistence with Covariates	103
3.9	Interquartile Difference in IEEs	113
A1	Hypothesis Tests with Covariates	124
A2	IEEs Across the Distribution (Joined Splines)	126

List of Figures

1	Reference-Dependent Utility Function with Loss Aversion	15
1.1	Savings over Sample Period	28
1.2	Distribution of Savings	29
1.3	Change SFS over Sample Period	32
1.4	Current SFS Versus Income Quintile	34
1.5	Change in SFS Versus Past Income Change	35
1.6	Expected SFS Versus Future Income Change	36
1.7	Indifference Curves with Reference Points	39
1.8	Savings Function with Loss Aversion	42
9	Indifference Curves with Reference Points with Uncertainty in Period 2 . . .	60
A1	Timeline of savings behaviour	64
2.1	GHQ (Likert) Distribution	73
2.2	Subjective Well-Being Over Sample Period	74
2.3	Loss Aversion with a Kink versus with a Jump	77
3.1	Age Distribution	91
3.2	Age Distribution	93
3.3	Parents' Life-Cycle Function	94
3.4	Parents' Earnings, Average and Age-Adjusted	95
3.5	Pattern of IEEs	104
3.6	Quantile Coefficients ($\hat{\theta}_h$)	111
3.7	Quantile Coefficients ($\hat{\theta}_l$)	112
3.8	Quantile Coefficients ($\hat{\theta}_l$ & $\hat{\theta}_h$)	113
A1	Robustness: Averaged Parental Earnings 2004-2008 ($\hat{\theta}_{l,q}$ & $\hat{\theta}_{h,q}$)	117
A2	Robustness: Averaged Parental Earnings 2004-2008 ($\hat{\theta}_{h,q}$)	118
A3	Robustness: Averaged Parental Earnings 2004-2008 ($\hat{\theta}_{l,q}$)	119
A4	Robustness: 60th Percentile Cutoff	120

A5	Robustness: 40th Percentile Cutoff	121
A6	Robustness: Exogenous Covariates	122
A7	Robustness: Endogenous Covariates	123
A8	Linear Splines	125
A9	Robustness: 3 Splines ($\hat{\theta}_{l,q}$)	127
A10	Robustness: 3 Splines ($\hat{\theta}_{m,q}$)	128
A11	Robustness: 3 Splines ($\hat{\theta}_{h,q}$)	129
A12	Robustness: Earnings Levels (All Splines)	130

Abstract

REFERENCE POINTS IN THE WILD: APPLICATIONS USING HOUSEHOLD DATA Cahal Moran

A thesis submitted to the University of Manchester
for the degree of Doctor of Philosophy, 2020

The aim of this thesis is to investigate the effect of reference points on households' decisions and outcomes using field data. The thesis looks at several different applications on multiple UK household datasets, all of which relate to the role of reference points in life cycle decisions about consumption, savings and income.

The first chapter studies the influence of reference points and loss aversion on consumption-savings behaviour. It has long been theorised that past income acts as a reference point, affecting individuals' consumption-smoothing. A simple two-period model shows that individuals below their reference point will save less, driven by a lower marginal propensity to save owing to loss aversion. The chapter proposes to measure reference points through subjective reports of financial changes in the British Household Panel Survey (BHPS), allowing individuals' own subjective perceptions to determine their reference points. Results show that *ceteris paribus* individuals below their reference point save less on average, which is driven by a marginal propensity to save about a third lower than those at or above their reference points. Further analysis investigates the effects of the SFS variables in more detail, including decomposing them using factor analysis and testing the influence of partner's SFS in couple households.

The second chapter tests the same class of models from another empirical perspective. Life cycle models featuring reference-dependent preferences and loss aversion imply that individuals who are below their reference point have lower overall utility but a higher marginal utility of consumption, which has important consequences for welfare and policy analysis. This chapter is the first to use 'experienced utility' in the form of subjective well-being (SWB) to corroborate this hypothesis. Two potential

determinants of reference points are proposed: last year's consumption (habit formation) and average regional consumption (relative income). Parameter estimates using the combined BHPS & Understanding Society dataset are consistent with the null hypothesis of loss aversion and indicate that it is characterised by a smooth utility function rather than a discontinuous one. The analysis favours a reference point based on habit formation as opposed to relative income, though the estimated parameters suggest model misspecification. Additional analysis decomposes the SWB measure into positive and negative affect to test the hypothesis that loss aversion is largely driven by negative affect as in psychological theories where losses are qualitatively distinct from gains, but this hypothesis is not supported.

The third chapter investigates intergenerational mobility in contemporary England and builds a model where parents' earnings act as an aspiration, modelled as a reference point, to explain the empirical results. Estimates of intergenerational mobility – Intergenerational Earnings Elasticities (IEEs) – have been shown to vary systematically across the joint earnings distribution of children and parents, but there is no agreed upon explanation for the observed pattern. A contemporary cohort dataset, the Longitudinal Study of Young People in England (LSYPE) is used to provide the most up to date estimates of IEEs in England, which are low on average though likely downward biased. Quantile regression is used to estimate IEEs across the distribution and results show they are lower for richer children but higher for richer parents. In order to explain this pattern a human capital model is constructed where parents' earnings act as an aspirational reference point for children. Children whose earnings are below their aspiration will have a higher marginal utility from investing in human capital and will compensate by increasing their future earnings faster than children who have reached their aspirational level.

Declaration

Some portion of the work referred to in this thesis has been submitted in support of an application for another degree at this university. Specifically, the description of the British Household Panel Survey/Understanding Society datasets, as well as some of the descriptions of the variables therein, were contained in my Master's dissertation at the University of Manchester.

Copyright

- i. The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the “Copyright”) and he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.
- ii. Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made **only** in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.
- iii. The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the “Intellectual Property”) and any reproductions of copyright works in the thesis, for example graphs and tables (“Reproductions”), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.
- iv. Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy (see <http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=487>), in any relevant Thesis restriction declarations deposited in the University Library, The University Library’s regulations (see <http://www.manchester.ac.uk/library/aboutus/regulations>) and in The University’s policy on presentation of Theses

Acknowledgements

Empirical analysis in this thesis uses data from the BHPS/Understanding Society study (Institute for Social and Economic Research, 2018; Institute for Social and Economic Research et al., 2019). Understanding Society is an initiative funded by the Economic and Social Research Council and various Government Departments, with scientific leadership by the Institute for Social and Economic Research, University of Essex, and survey delivery by NatCen Social Research and Kantar Public. The research data are distributed by the UK Data Service.

I am also grateful to the Centre for Longitudinal Studies (CLS), UCL Institute of Education, for the use of the Longitudinal Study of Young Persons Earnings (LSYPE) data and to the UK Data Service for making them available. However, neither CLS nor the UK Data Service bear any responsibility for the analysis or interpretation of these data.

I would like to thank my two PhD supervisors James Banks and Horst Zank for their supervision, as well as faculty members Agnes Kovacs, Ralf Becker, Martyn Andrews, Peter Backus, and Simon Peters for help when I was struggling with a problem. Among the PhD students I would like to thank Atiyeh, Esmael, Laura, Yan, Keila, Tad, Myroslav, Nico, Alex and Thien, for providing social, emotional and, on occasion, academic support. I would like to thank my family and friends for raising and accompanying me (respectively), though there are too many names to list in full. However my mum and dad Rachel and Richard, as well as all of Good Friends, deserve special mention. Finally, I wish to thank my recently passed colleague James Lincoln, who taught me half of what I know about coding and who will be sorely missed by all of us.

Introduction

This PhD thesis measures the causes and consequences of reference points for household behaviour with the aim of contributing to the field literature on prospect theory. I aim to understand how reference points are formed in the real world and how they influence economic outcomes — specifically consumption smoothing, subjective well-being, and intergenerational mobility. I develop techniques to assess the effects of reference points in standard household data, an area which has not received much attention in the literature. Reference points often have unique implications for understanding individual behaviour: two individuals who are otherwise similar may behave differently simply because of their reference points, and will also have different levels of welfare. The policy implications could entail modifying reference points directly, or simply accounting for the differences in behaviour when designing programs and incentives. However, this thesis is relatively silent on explicit policy and welfare analysis, instead aiming to establish the empirical implications of reference points for behaviour.

The thesis consists of three chapters, all of which take a different approach to measuring the effects of reference points in different contexts. The first two chapters use the same dataset, the British Household Panel Survey (BHPS) (Institute for Social and Economic Research, 2018) and Understanding Society (USoc) (Institute for Social and Economic Research et al., 2019), though the first chapter only uses the BHPS part due to unavailability of a key variable in USoc. Both chapters are testing similar mechanisms – the effect reference points have on optimal consumption-savings decisions in a standard life-cycle framework – but the first investigates short-term consumption-smoothing behaviour while the second uses subjective well-being to measure the marginal utility of consumption. In other words, the first uses revealed preferences while the second uses stated preferences to measure the effect of reference points. The third and final chapter uses a different dataset, the Longitudinal Study of


Young People in England (LSYPE) (Centre for Longitudinal Studies, 2018), to investigate intergenerational mobility in England. It explains the observed pattern of mobility across the distribution in a model where parents' earnings act as an aspirational reference point for children.

The motivation for this thesis was to apply insights from prospect theory to large scale datasets. The seminal paper of Kahneman and Tversky (1979) was, at the time of writing, the second most cited paper in economics, yet empirical applications of prospect theory in the real world (as opposed to the laboratory) are surprisingly rare. One of the key innovations of prospect theory over standard expected utility theory was the idea that individuals compare their level of consumption to a reference point. Furthermore, individuals experience 'loss aversion': they value a 'loss' (being below their reference point) as higher in magnitude than an equivalent 'gain' (being above their reference point). It has become increasingly accepted among researchers that individuals compare themselves to reference points in many contexts and that this affects their economic decisions. Yet the issue of what determines reference points in different contexts, and the quantitative effect reference-dependent preferences have on behaviour, is still not fully understood (Camerer, 2000; Starmer, 2000; Schmidt et al., 2008; Shleifer, 2012; Barberis, 2013).

Generically reference-dependent utility functions take the form $v(c_i, r_i)$ where c_i is consumption and r_i is a reference point to which the individual compares their level of consumption. Figure 1 shows such a function. It is assumed that individuals are loss averse so that there is a utility penalty from being below one's reference point, creating a 'kink' when $c_i = r_i$ so that $v(r_i, r_i) = 0$. Formally $v'(x) < v'(-x) \quad \forall x$: for a given $|c_i - r_i|$, when $c_i < r_i$ the slope of the utility function is steeper than when $c_i \geq r_i$. It follows that welfare losses when the individual is below r_i are higher than with standard preferences: due to this kink, the slope of the utility function becomes steeper so that marginal utility below r_i is higher.

This thesis follows the Koszegi and Rabin (2006, 2007, 2009) specification in making the degree of loss aversion linear for tractability and for clarity of predictions. The exact form of this utility function will be specified in each chapter, adapted for each specific purpose. A standard graphical representation of reference dependent utility functions with loss aversion is nevertheless useful to consider and is presented in Figure 1. The general message is that the kink in utility at r_i means that individuals will behave differently depending on where their consumption is relative to their reference point. In Chapter 1, individuals who are below their reference point will save less,

Figure 1: Reference-Dependent Utility Function with Loss Aversion



A utility function featuring reference-dependent preferences and loss aversion. Above the reference point r_i the utility function is a standard concave function. Below r_i there is a ‘kink’ and the slope of the utility function becomes steeper so that the marginal utility below r_i is higher.

which follows naturally from their higher marginal utility of consumption – and this aspect of the theory is tested directly in Chapter 2, using ‘experienced utility’ in the form of subjective well-being. In Chapter 3, children will accrue human capital to increase their earnings at a faster rate when they are below r_i for the same reason: their marginal utility of an extra unit of earnings is higher.

The task throughout this thesis is to use these theoretical insights to explain what we observe in the data and therefore show that reference points have important implications for behaviour and outcomes. However, in order to do so we must know what determines r_i , an issue which is plagued by theoretical and empirical problems. There can be no doubt that r_i varies by context and so can be manipulated (at least to a degree) by the individual’s environment. There have been several approaches taken in the literature which take advantage of this point to study reference points in particular contexts.

In laboratory tests researchers can construct the environment to create a salient reference point to affect subjects’ behaviour. Kahneman and Tversky (1979)’s original laboratory evidence showed the status quo acts as a reference point and consequently that individuals care about changes in wealth as well as wealth levels. Disparities between Willingness-to-Pay and Willingness-to-Accept measures generally show that when given an item at the start of an experiment, individuals demand more money as payment for it than they are willing to pay when they do not have it, consistent with the value of the item acting as a reference point and subjects being averse to ‘losing’

the item (Knetsch and Sinden, 1984; Boyce et al., 1992; Bateman et al., 1997). This has been christened the endowment effect by behavioural economists. Other experiments ask subjects to choose between hypothetical gambles and use non-parametric methods to estimate the value function directly (Gonzalez, 1999; Abdellaoui, 2000), using the status quo as a reference point, and these generally confirm the existence of loss aversion.

Though laboratory tests are useful for precisely illustrating key cases where individual behaviour follows reference-dependent models rather than more standard models, there are questions about their external validity. As such there are an increasing number of field applications of prospect theory and these typically use cases where the reference point is *naturally* salient. One example is zero net winnings per day for gamblers (Ali, 1977) where gamblers take more and more risks to break even. Another prominent example is the notion that New York taxi drivers stop working after they reach a target income for the day (Camerer et al., 1997; Farber, 2005, 2008; Crawford and Meng, 2011). Individuals also withdraw from their pensions consistent with a socially normed retirement age for labour supply (Seibold, 2017), and there is evidence that unemployment benefits act as a reference point, affecting reservations wages in job search (Dellavigna et al., 2017). All of these are cases which are useful for understanding behaviour in real world contexts and advising policy.

It is harder to measure reference points in more general settings such as households' consumption-savings decisions, since they are subjective to each individual and could be influenced by a variety of factors including peers, the status quo, and both past and future expectations (Barberis, 2013). In many cases it will not be possible for the researcher to impose a specific reference point so available data and methods must be used to infer individuals who are likely to be above versus below their reference points. Their behaviour can then be investigated to see whether or not it is consistent with the basic logic of Figure 1. It is the aim of this thesis to shed some light on this final case, where research is less plentiful.

A recent line of literature has questioned whether loss aversion is as universal as its proponents sometimes claim. Gal and Derek D. Rucker (2018) review cases which are frequently cited as evidence for loss aversion and show repeatedly that alternative explanations are possible. For example, the endowment affect above could instead be explained by the inertia of potential sellers, and List (2004) famously showed that it disappears under market conditions i.e. if traders can gain experience from repeated trading. Although loss aversion has been used to explain the equity premium puzzle

(Benartzi and Thaler, 1995) – the implausibly high level of risk aversion implied by the high returns on stocks relative to bonds – there are a number of alternative explanations, for instance habit formation (Constantinides, 1990).

Yechiam (2018) claims that much of the original laboratory evidence for loss aversion was interpreted selectively, and in fact some of the studies cited by the original Kahneman and Tversky (1979) paper do not support the notion of loss aversion. Brooks et al. (2014) show that in an experiment specially designed to test the distinct components of prospect theory, loss aversion is not always observed. It is therefore my view that loss aversion is not universal and since it drives many of the theoretical predictions in this thesis, special attention is paid to whether alternative explanations for the empirical results are possible in the first and third chapters (it is less obvious what the alternative would be in the second chapter). In no case can I claim that the discussion is exhaustive, but in each case the most prominent and feasible alternatives are ruled out. Each chapter of the thesis will now be summarised in detail.

The first chapter illustrates that reference-dependent preferences and loss aversion drive intertemporal consumption-smoothing behaviour. The idea is that when individuals experience a financial change their past situation acts as a reference point which makes them slow to adjust, affecting their consumption-smoothing behaviour. This insight originated with the relative income hypothesis of James Duesenberry (1949) but has been neglected in the contemporary consumption-savings literature, which typically uses the permanent income hypothesis (Jappelli and Pistaferri, 2010; Attanasio and Weber, 2010). As reference points are inherently subjective I propose using Subjective Financial Situation (SFS) variables, which are available in the BHPS, to measure them. I show that there is a disparity between self-reported SFS and observed income, suggesting the SFS variables give a window into the perceptions of individuals which objective variables cannot.

There are three SFS variables: current, expected and change in SFS, and I interpret the latter as a measure of an individual's position relative to their reference point, a key assumption which is discussed in more detail in the chapter. The intuition is that if two individuals have the same current financial situation but one has experienced a positive change while the other has experienced a negative change, we can plausibly infer that the latter is below their reference point. I show in a two-period model of consumption-smoothing with reference-dependent preferences and loss aversion that individuals below their reference point will have both a lower level of savings owing to a lower marginal propensity to save (MPS), because they smooth consumption into

the current period to reduce their losses.

Regressions are run of self-reported individual savings on income, an indicator for whether an individual is below their reference point (as measured by a decline in SFS), and an interaction between the two. Whether or not the MPS exhibits loss aversion can be tested by whether the interaction term is negative and statistically distinct from 0. The data reject the null hypotheses of no loss aversion against the one-sided alternative of loss aversion at all conventional levels of significance. Results show that the MPS of individuals below their reference point is about a third lower than the MPS of those who are above it.

Using the lagged expectations SFS variable, the chapter investigates whether unexpected changes can also be considered a measure of reference points. If an individual expected things to improve last year but they did not, they can be said to be below their reference point. Regression results show that having SFS below one's past expectations reduces the MPS by about a third. This implies that reference points depend on past expectations as well as realised changes. Further analysis is less prescriptive and uses all three variables (current, change in and expected SFS) in a factor decomposition, and regressions show the two extracted factors influence saving. These results do not speak as directly to reference-dependent preferences and loss aversion as the previous results, but nor are they inconsistent with them. Finally, the sample is restricted to couple households to test whether one's partner's SFS influences one's own savings, but there is no effect.

The second chapter makes use of 'experienced utility' in the form of subjective well-being (SWB), as well as data on food expenditure, to estimate the parameters of a utility function featuring reference-dependent preferences and loss aversion. This is to my knowledge the first attempt to test this aspect of prospect theory using SWB, with the most similar paper being Finkelstein et al. (2012) who do the same for state-dependence in health status. In the harmonised BHPS/USoc dataset I use two potential determinants of reference points: last year's consumption similar to habit formation (Constantinides, 1990), and average regional consumption similar to the relative income hypothesis (Duesenberry, 1949). I use the GHQ-12 questionnaire, which was designed as an indicator of mental health but has increasingly been used as a measure of SWB in the literature (Clark and Oswald, 1995; Bayliss et al., 2014; Jones and Wildman, 2008; Oswald, 1997).

This chapter tests for reference-dependence and loss aversion through direct estimation of the model parameters using a regression model. For the habit formation

measure of reference points, the null hypothesis of no loss aversion is tested against the one-sided hypothesis of loss aversion and cannot be rejected at conventional significance levels. However, other parameter estimates indicate either some misspecification of the model or insufficient variation in consumption to identify the key effects. For the relative income measure of reference points, the null hypothesis of no loss aversion is tested against the one-sided hypothesis of loss aversion and is rejected at conventional significance levels, indicating no reference-dependence in terms of relative income.

Following Bayliss et al. (2014) I decompose GHQ into positive and negative affect and show that consumption is driven more by negative affect but loss aversion is not, which is inconsistent with the idea proposed in the literature that losses are psychologically distinct from gains (Argyle, 1987). Instead, the results show that negative utility is more affected by food consumption than positive utility. Overall, I conclude that the evidence for reference-dependent preferences and loss aversion is mixed, but that this type of data is a promising way to estimate ‘behavioural’ utility functions in the future.

The third chapter measures the degree of intergenerational mobility across the joint distribution of parents’ and children’s earnings in contemporary English cohort data, and explains the pattern of results as an implication of parents’ earnings acting as an aspiration for children. I estimate the effect of parents’ on children’s earnings – the Intergenerational Earnings Elasticity (IEE) – using the Longitudinal Study of Young People in England (LSYPE), the most recent study available in the country. Results show that the aggregate IEE is low, though this is likely downwards biased. I then estimate IEEs across the distribution by using quantile regression of children’s earnings on linear bins of parents’ earnings, a method which has been used in the literature but for which there is no clear theoretical interpretation of the results (Grawe, 2004).

The quantile approach shows that conditional on a bin of parents’ earnings, the IEE falls as the quantile of children’s earnings rises: richer children are associated with lower IEEs. Conversely, conditional on a quantile of children’s earnings the IEE is higher for upper bins of parents’ earnings: richer parents are associated with higher IEEs. Other variables also have different impacts across the distribution, with ‘human capital’ variables such as education, occupation and health compressing the distribution while non-control variables such as race and gender display more ambiguous patterns. However, none of these covariates fully explain the effect of parents’ earnings on childrens’, nor do they eliminate the observed pattern across the different IEEs.

This result is inconsistent with standard explanations for earnings persistence such

as intergenerational altruism or credit constraints which have been modelled by previous studies (Han and Mulligan, 2001; Grawe, 2004). In order to explain it I construct a model where parents' earnings act as an aspiration for children. Children whose earnings are below their aspiration will have a higher marginal utility of investment in human capital and will compensate by increasing their earnings faster than children who have reached their aspirational level. Unlike standard models of intergenerational mobility this implies the IEE will depend on the relative distance between parents' and children's earnings, rising when children's earnings are lower than parents' and falling when the converse is true — consistent with my empirical results.

This has the twin implications that the IEEs will be lower for upper quantiles of children's earnings, and higher for upper bins of parents' earnings, and both hypotheses can be tested formally. Interquartile regression results show that for a given bin of parents' earnings, the null hypothesis that the IEE is equal across quantiles can be rejected against the one-sided alternative that it is higher for lower quantiles. Furthermore, results show that for a given quantile, the null hypothesis that the IEEs are equal for both bins of parents' earnings cannot be rejected against the one-sided alternative that the IEE for the upper bin is higher. In summary, both hypotheses of the model are consistent with the data and this result survives a number of robustness checks.

The third chapter is the most indirect approach to estimating the effect of reference points out of the three: the model is argued to produce a pattern consistent with the results even though we cannot directly observe reference points or otherwise test the key mechanisms in the model. Subsequently there is more uncertainty about whether the reference point framework is necessary to generate the results or whether a more standard mechanism would do, though as mentioned above several alternative explanations are ruled out. The purpose of the chapter is to suggest an interpretation of a method which has been used but not fully interpreted in the literature, to show that it is a credible explanation in this particular context, and to propose it as a general technique for detecting the influence of reference points when we cannot observe them. Future research will hopefully develop and refine this approach further, as well as testing other implications of the model itself.

Chapter 1

Subjective Financial Situation as Reference Points

1.1 Introduction

This chapter argues that reference points and loss aversion affect short-term consumption-smoothing behaviour by reducing the marginal propensity to save of individuals below their reference points. It proposes to measure reference points using Subjective Financial Situation (SFS) variables, allowing individuals' own subjective reports to determine their reference points. It is shown using British panel data that individuals who are below their reference points have a marginal propensity to save about a third lower than those at or above their reference points. This result is consistent with a two-period savings model featuring reference-dependent preferences and loss aversion. Calculations show that if everyone in the sample who is currently below their reference point were instead at or above it, the savings rate would increase by two percentage points, a large macroeconomic effect.

The main contributions of this chapter are to propose the SFS variables as an empirical measure of reference points, and show that savings functions exhibit the property of loss aversion. In doing so the chapter contributes to three literatures. Firstly, it contributes to the literature measuring the effects of reference points in field settings, which lacks measures of reference points in standard household data (Starmer, 2000; Shleifer, 2012; Barberis, 2013). Secondly, it contributes to the life cycle literature on the effect of income changes, which does not typically incorporate loss aversion (Jappelli and Pistaferri, 2010; Attanasio and Weber, 2010). Thirdly, it contributes to the older literature on the relative income hypothesis (Keynes, 1936; Duesenberry, 1949),

which envisioned similar mechanisms but did not test them using sophisticated contemporary data and methods.

The intuition behind the result in this chapter is adaptation: individuals acclimatise to a given level of income which acts as a reference point, affecting their response to short-term fluctuations. If an individual is used to an income of £2000 per month it will be difficult for them to adjust to a £500 fall, so they are likely to borrow more than another individual who has earned £1500 a month for years. Such behaviour will also be affected by subjective perceptions: two individuals may both receive a £500 raise, but one may have anticipated it while another didn't; or one may be surrounded by peers who have received £1000 raises while the other is surrounded by peers who lost their jobs. This will affect the individuals' reference points and resultant consumption-smoothing behaviours.

Although a lagged effect of income on consumption and savings is a well known empirical regularity and a common implication of standard models, it is not clear whether this is a result of reference-dependence and loss aversion. This chapter devises a unique test for the latter hypothesis by testing for a difference in the marginal propensity to save between those above versus those below their reference points. It shows that savings functions are characterised by loss aversion, suggesting the result is not driven by other potential explanations for the effect of income changes in the life cycle literature (Jappelli and Pistaferri, 2010; Attanasio and Weber, 2010).

The three Subjective Financial Situation (SFS) variables in the chosen dataset, the British Household Panel Survey (BHPS), pertain to current financial situation, expected change next year, and realised change since last year. Empirical comparisons show that experienced income changes correlate with reported changes in SFS as would be expected, but that there is substantial heterogeneity in SFS within each type of income change. For example, some who have experienced an increase in reported income report a decline in SFS and vice-versa. Analogous results are shown for the current and expectation SFS variables. It is concluded that SFS captures things missed by conventional variables such as income. This includes other aspects of financial situation such as consumption needs, assets, debt and unreported income streams, but also includes subjective perceptions and changes in circumstance.

Regression of self-reported savings on the SFS variables and income allow estimation of the marginal propensity to save (MPS) for those at or above, versus those below their reference points. Results show that on average, those below their reference point save less than those at their reference point, equivalent to about £7,000 less income.

This is driven by the fact that the MPS of individuals below their reference point is about 40% higher than for those at or above it. The latter result confirms that the effect of income changes is best characterised using loss aversion among those below their reference point, who save a lower proportion of their income in order to reduce their perceived loss in consumption.

Additional analysis tests whether unfulfilled expectations are also a determinant of reference points: if an individual's expectations SFS in the previous year exceeded their realised change, they can be said to be disappointed and therefore below their reference point. Using this measure, the predictions of the model are also borne out, suggesting reference points depend on past expectations as well as realised changes. Effects are similar in magnitude to the results using the raw change SFS summarised above.

Further results use factor analysis to decompose the three SFS variables into two underlying processes to 'let the data speak for itself', rather than assuming reference points are best measured by the change SFS only. The first factor is interpreted as reporting style, consistent with SFS variables capturing subjective perceptions missed by conventional variables. The second factor is a separate component which measures negative financial situation, consistent with the unique influence of losses implied by loss aversion and hypothesised in the psychology literature (Diener, 1984; Argyle, 1987). It is shown that the first factor has a positive effect on savings and the second has a negative effect on savings, which is consistent with reference-dependent preferences and loss aversion but does not speak as directly to them as earlier results.

It is possible that Subjective Financial Situation are household rather than individual-level phenomena, especially for couple households who are likely to pool their finances. To investigate this the sample is restricted to the heads of these households – along with their partner's SFS – and within-household differences in SFS are shown to be common, even among households who report sharing resources equally. This supports the notion that SFS is partly driven by subjective perceptions, since it shows that two people can have distinct evaluations of one financial situation. Despite these differences, regressions including both the head's and the partner's SFS factors show that partner's SFS does not predict savings. It is concluded that only one's own subjective perceptions determine savings behaviour. Appendix A.2 runs two robustness checks and finds them to be generally consistent with the main analysis.

The chapter proceeds as follows. Section 1.2 reviews the literature on the determinants of reference points with special reference to case of consumption and savings. Section 1.3 outlines the BHPS and explores the SFS variables in more detail. Section 1.4 builds the theoretical model and summarises its predictions, while Section 1.5 discusses the associated econometric methodology used to estimate savings functions. Section 1.6 presents the empirical results, including the additional analysis with factor decomposition and within-household comparisons. Section 1.7 concludes and discusses implications for future work.

1.2 Literature Review

A long-standing literature holds that perceptions of past income act as a reference point in consumption and savings decisions, causing a short-term ‘adaptation effect’ as individuals adjust to changes. This was a core part of Duesenberry (1949)’s relative income hypothesis and can even be found in Keynes (1936). It is also one of the key innovations of prospect theory (Kahneman and Tversky, 1979): the idea that utility levels depend not on final wealth positions but positions relative to a reference point. This chapter contributes to this literature by providing an empirical measure of the reference points which create the adaptation effect and showing that it is driven by loss aversion.

Modern consumption-savings research has moved away from adaptation for both theoretical and empirical reasons. Modigliani and Brumberg (1955) formalised some key insights of Duesenberry (1949) but took a different view of the relative standard of income, interpreting it as an individual’s “permanent income”, which resulted in what became known as the Life Cycle Model (LCM). This built both long term and short-term savings behaviour into a single framework where individuals smooth their marginal utility across periods. The literature then went on to distinguish between different types of income changes such as transitory versus permanent; expected versus unexpected; and aggregate versus idiosyncratic (Jappelli and Pistaferri, 2010; Attanasio and Weber, 2010). While useful distinctions, these seek to answer questions which can be measured using conventional ‘objective’ variables like income, whereas reference points depend on subjective variables.

It is common to use habit formation as an implicit reference point in the LCM (Constantinides, 1990; Chetty and Szeidl, 2009). This results in consumers reacting slowly

to income changes because they have developed a stronger preference for commodities they have consumed in the past. Bowman et al. (1999) used an earlier version of prospect theory with a reference point partially dependent on consumption. However, these models can result in the individual ‘starving’ themselves, accumulating wealth and keeping consumption low in early periods so that utility will be higher in later ones. As shown by Michaelides (2002), habit formation models of this style imply an unreasonable time series for both wealth and consumption, suggesting individuals’ past consumption is not an appropriate choice for their reference points.

Koszegi and Rabin (2006, 2007, 2009) develop a version of prospect theory in which reference points are formed endogenously as a result of rational expectations over uncertain future states. In the case of intertemporal savings behaviour these expectations create an asymmetry where individuals tend to consume current wealth gains and postpone future wealth losses, since they will have adjusted their reference point downwards in the future (Koszegi and Rabin, 2009). Pagel (2016, 2018) applies the rational expectations framework to the LCM. Pagel (2016) shows that this model can reconcile a number of stylised facts in the savings literature, including excess smoothness and sensitivity (Campbell and Deaton, 1989), a hump-shaped consumption profile over individuals’ lifetimes, and a drop in consumption at retirement (Banks et al., 1998; Browning and Crossley, 2001). Pagel (2018) extends the framework to portfolio choice, showing that the ‘news-utility’ created by expectations may cause investors to ignore news and rebalance their portfolio less than is required.

Although Pagel’s models have strong predictive power when it comes to the macroeconomic consumption and savings data, Masatlioglu and Raymond (2016) have shown that the rational expectations reference point has a number of theoretical shortcomings. Firstly, it can be expressed as a version of Rank Dependent Utility (Quiggin, 1982), so that it is not clear whether the individual is ‘loss averse’ or ‘pessimistic’ — therefore Pagel’s model cannot be considered a test of reference-dependent preferences. Secondly, certain specifications can plausibly violate first order stochastic dominance, an issue which was also present in the original version of prospect theory (ibid). Thirdly, the most tractable version of the model cannot avoid the criticism made by Rabin (2000) of expected utility theory: that preferences over small stakes gambles imply implausible behaviour over large stakes gambles.

I contribute to the reference point literature by providing an empirical measure of reference points in the context of short-term consumption-smoothing behaviour. Unlike most field studies in the reference point literature the SFS measure is available

in many household datasets and could be used in a variety of applications. SFS measures subjective perception of financial changes, avoiding reference points which are unfeasibly complex functions of both preferences and the environment. The model is similarly simple, leaving the reference point as an exogenous variable to be measured empirically, and predicts that individuals below their reference point save a lower proportion of their income¹.

1.3 Data and SFS Variables

This section describes the dataset used in this chapter, the British Household Panel Survey (BHPS) (Institute for Social and Economic Research, 2018). The BHPS is an annual survey of over 5,000 randomly selected households (around 10,000 individuals) who, where possible, were interviewed every year between 1991-2008. Only adults – defined as age 16 or over – were interviewed for the survey, which made use of a combination of face-to-face interviews, telephone interviews, proxy answers and self-completed questionnaires, with some use of imputed values for variables such as income (Taylor et al., 2010). There are a total of 240,000 person-years in the dataset, with 33,000 people and an average panel length of 6.5 years. Table 1.1 summarises descriptive statistics for in the sample other than the main variables of interest, which will be discussed below.

The BHPS contains a variable for self-reported savings where each individual is asked “About how much on average do you manage to save a month?” Savings are censored at zero, which means that individuals who report zero could be saving exactly 0 or could be borrowing i.e. ‘saving’ a negative amount. 82,000 person-years report positive savings while 137,000 report not saving, and the remaining observations are coded as genuinely missing. When savings are annualised and deflated to 2005 price levels the average amount of savings per year is £680, including zeros. The average saved for those with positive savings is £2000. Figure 1.1 shows the average

¹The model draws heavily from the unpublished paper by Cohen (2015), including in the logic of his proof. The main difference is that Cohen’s reference point is equal to income for the current period where here it is dependent on income for the previous period. Since Cohen’s reference point is for the current period there are no obvious implications of his model for the response to short-term savings behaviour, whereas this chapter clearly predicts the adaptation effect and what drives it. Furthermore, in this chapter the reference point is exogenous, to be measured empirically with individual-level data, whereas in Cohen’s model the reference point is of interest only insofar as it drives the macroeconomic ‘stylised facts’ his model seeks to explain.


Table 1.1: Economic and demographic variables in the BHPS

	Mean	Standard Deviation
Panel A: Standard Variables		
Age	45.28	18.65
Number of persons in household	2.80	1.45
Number of own children in household	0.51	0.93
Real income (£)	13,395	10,572
Hours worked (full sample)	17.6	18.8
Hours worked (employed)	33.5	12.1
Subjective health status (1-5 scale)	2.18	0.95
Panel B: Dummy Variables (0-1)		
Male	0.46	0.50
White	0.92	0.26
Household		
Own/Mortgage house	0.70	0.46
Rent house	0.25	0.43
Marital Status		
Living as couple	0.11	0.31
Married	0.53	0.50
Widowed	0.08	0.27
Divorced	0.05	0.22
Qualifications		
Degree or higher	0.17	0.38
Some qualifications	0.60	0.49
No qualifications	0.23	0.42
Employment Status		
Self-employed	0.07	0.25
Employed	0.50	0.50
Unemployed	0.04	0.19
Retired	0.20	0.40
Student	0.07	0.26
Long term sick/disabled	0.06	0.24
Region		
London & Southeast	0.17	0.37
North	0.19	0.39
Midlands & East	0.19	0.39
Wales & Southwest	0.20	0.40
Scotland & Channel Islands	0.15	0.36
Northern Ireland	0.09	0.29

The means and standard deviations of variables in the BHPS from 1991-2008, measuring economic and demographic characteristics for the entire sample ($N = 238,996$). Panel A variables are all positive real numbers, while Panel B variables are measured in proportions between 0 and 1. All given moments exclude missing values for that variable.

real savings over the sample period, which rises over time — even during the recession, something which has previously been documented as ‘saving on a rainy day’ (Alan et al., 2012). Figure 1.2 shows the distribution of savings for those with positive savings, with a peak at £200 – 300 and another peak at £600 – 700. In general the proportion of individuals saving in a particular bin declines steadily as the amount of savings increases, though the histogram is cut off at £3000.


Figure 1.1: Savings over Sample Period



The evolution of self-reported savings in the BHPS over time. Each individual is asked “About how much on average do you manage to save a month?” which is naturally censored at zero, with 30% of the sample reporting positive savings (this graph includes zeroes in the average).

Typically when economists measure savings they derive it financially — as income minus consumption, or a change in assets. The savings measure in the BHPS is distinct from these measures and could be termed ‘active’ saving: deliberately moving resources into a different account or asset. This is analogous to the ‘discretionary saving’ measure of Borsch-Supan et al. (2003) but it is even more deliberate, and was used in a test of precautionary savings in the BHPS by Guariglia (1998). Unfortunately a financially derived variable for savings is not possible in the BHPS, but the active savings measure is advantageous in some respects for this chapter. The survey question

Figure 1.2: Distribution of Savings



The distribution of the self-reported savings variable in the BHPS, with zeros excluded and a cutoff at £3000.

abstracts from long term decisions such as pensions, housing, and bills so it is relevant for individuals actively ‘putting money away’ to smooth consumption as their income fluctuates over the short term. Furthermore, this chapter is focused on subjective perceptions and as such it is important to know what people believe they have saved – a financially derived magnitude may not be salient to the individual.

The three Subjective Financial Situation (SFS) variables in the BHPS ask each individual to evaluate their own financial situation. These variables relate to the individual’s financial situation this year, expected financial situation next year, and change in financial situation since last year. Respondents only give ordinal answers to these questions, but as the model developed later illustrates, this will be all that is needed to test the predictions of the model. An additional variable asks those who have experienced a change in financial situation since last year the reason their financial situation changed, which will be useful to control for omitted variable bias in the regression analysis later.

Table 1.2: Current Financial Situation Summary Statistics

Current SFS	Proportion in sample	Mean Savings (£)
Living comfortably	30.12%	1389.77
Doing alright	36.49%	655.91
Just about getting by	25.36%	221.27
Finding it quite difficult	5.60%	105.96
Finding it very difficult	2.42%	42.86

Possible answers to the question “How well would you say you yourself are managing financially these days? Would you say you are...”. It also shows the proportion of respondents who gave each answer, as well as their average reported real savings (including zeroes). The sample is restricted to cases where no SFS variables are missing (219,629 observations).

Table 1.3: Expected Financial Situation Summary Statistics

Expected Financial Situation	Proportion in sample	Mean Savings (£)
Better	27.67%	770.65
About the Same	61.18%	699.02
Worse	11.16%	719.19

Possible answers to the question “Looking ahead, how do you think you yourself will be financially a year from now, will you be...”. It also shows the proportion of respondents who gave each answer, as well as their average reported real savings (including zeroes). The sample is restricted to cases where no SFS variables are missing (219,629 observations).

The details of the three SFS variables and summary statistics are given in Tables 1.2, 1.3 and 1.4, including the average level of savings for those who report each answer (the tables show only observations which contain all three SFS variables available, about 220,000 person-years or 91% of the sample). As would be expected, in the raw data reporting a better current financial situation is associated with higher savings, while those who have experienced a positive change have higher savings than those who have experienced a negative change. The correlations for the expectations variable are less straightforward, with savings higher for both individuals who expect a positive change and for individuals who expect a negative change — relative to those who expect roughly no change. This is likely because positive expectations are associated with a positive experienced change since people have myopic expectations and expect changes to continue, a conjecture supported by the cross-correlations in Table 1.5.

Figure 1.3 shows the answers to the change SFS variable from 1991-2008 and they follow a plausible pattern given what is known about the UK economy during this period. Over the 1990s the proportion of individuals whose financial situation has either

Table 1.4: Change in Financial Situation Summary Statistics

Change in Financial Situation	Proportion in Sample	Mean Savings (£)
Better	26.95%	1172.32
About the Same	50.19%	618.53
Worse	22.86%	414.36

Possible answers to the question “Would you say that you yourself are better off, worse off or about the same financially than you were a year ago?”. Shows the proportion of respondents who gave each answer, as well as their average reported real savings (including zeroes). The sample is restricted to cases where no SFS variables are missing (219,629 observations).

Table 1.5: Change in Financial Situation Summary Statistics

Correlation	Current SFS	Change SFS	Expected SFS
Current SFS	1		
Change SFS	0.3494	1	
Expected SFS	-0.0133	0.1882	1


Correlations between the three subjective financial situation (SFS) variables with the sample restricted to cases where none are missing (219,629 observations).

improved or stayed the same since the previous year increases, with a concomitant fall in the proportion of those whose has declined. During the 2000s there is a slight decline in those whose situation has improved and a flatlining of those whose has declined. There is also a continuation in the upward trend of those who report roughly no change. However, when the financial crisis hits in 2008 there is a huge increase in the proportion of individuals reporting their financial situation has gotten worse and a decrease in the proportion reporting that it has got better or stayed the same.

When it comes to measuring financial circumstances there are two main differences between the SFS variables and objective variables such as income. Firstly, individuals may have unreported assets and liabilities which will enter into their considerations but are not generally asked about in surveys. As the SFS variables ask for an assessment of individual’s financial situation as a whole, these will be captured better by the SFS variables. Secondly, the SFS variables allow for differences in perceptions between two individuals who face the same objective financial situation. As discussed in the introduction there is likely to be considerable variation in how each individual perceives a given change.

Figure 1.4 shows how the proportion who give each answer to the current SFS question (within the bars) relates to income quintiles. There is a positive correlation

Figure 1.3: Change SFS over Sample Period



The three possible answers to the change SFS variable over the sample period. Those whose situation got ‘better’ than last year are given by the solid line, those whose situation stayed ‘about the same’ are given by the dot-dashed line, and those whose situation got ‘worse’ are given by the dashed line.

between giving a top answer (or not giving a bottom answer) and having more income. Despite this, every income quintile contains people who give each SFS answer: even at the top income quintile about 5% answer that they are finding it quite difficult or finding it very difficult. Figure 1.5 and Figure 1.6 do the same thing for the change SFS and expected SFS variables respectively, splitting the sample into those who have experienced (or go on to experience) income increases, decreases or small changes². Once more, each income category contains a substantial number giving each SFS answer, illustrating that perceptions can differ markedly from objective measures.

The additional SFS variable asks the reason an individual’s situation changed among those who report any change, and is reported for all but the first two waves of the BHPS. Selected answers are shown in Table 1.6 and most importantly for our purposes, the answers ‘good financial management’ and ‘less savings’ can be used as controls to account for any reverse causality in the later estimation. The table shows that the proportion giving each answer is low — around 3% of those who report a change. Its inclusion/exclusion does not affect any of the results reported in this chapter.


²Defined as under £1,000 in absolute magnitude.

Table 1.6: Reason Financial Situation Changed

Answers	Proportion in sample	If improved	If worsened
Earnings increased	14.64 %	14.64 %	-
Benefits increased	1.88%	1.88%	-
Investment income in- creased	0.60%	0.60%	-
Less expenses	3.15%	3.15%	-
Windfall payment	0.67%	0.67%	-
Earnings decreased	5.53%	-	5.53%
Benefits reduced	0.69%	-	0.69%
Investment income de- creased	0.60%	-	0.60%
More expenses	10.03%	-	10.03%
One off expenditure	0.43%	-	0.43%
Good financial man- agement	1.29 %	1.29 %	-
Less Savings	0.27%	0.27%	-
No change in income/- expenses	0.04%	0.04%	-
Better off: other reason	2.25%	2.25%	-
Worse off: other reason	2.20%	-	2.20%
Another reason	0.01%	0.01%	-
No change	55.39 %	-	-


Answers to the follow up question “Why is that?”, asked to individuals who report they are either financially better or financially worse off than last year. This variable is only measured from 1993-2008 i.e. it is missing for the first two waves.

Figure 1.4: Current SFS Versus Income Quintile




This graph shows the current SFS variable on the y-axis with income quintile on the x-axis. Each colour represents the proportion of the sample in a particular income quintile who give an answer to the current SFS question. Although there is a positive correlation between SFS and income, each income quintile contains every SFS answer.

Figure 1.5: Change in SFS Versus Past Income Change



This graph shows the change SFS variable on the y-axis with the experienced change in income since last year on the x-axis. Each shade represents the proportion of the sample who have experienced a decline, increase or roughly no change (defined as under £1000 in magnitude) who give a particular answer to the change SFS question.

Figure 1.6: Expected SFS Versus Future Income Change



This graph shows the expected SFS variable on the y-axis with their realised income change in the following year on the x-axis. Each colour represents the proportion of the sample who go on to experience a decline, increase or roughly no change (defined as under £1000 in magnitude) who give a particular answer to the expected SFS question.

1.4 Model

This section develops a model of intertemporal consumption and savings choices which features reference-dependent preferences and loss aversion. The model is a generalisation of Cohen (2015)'s model, where the reference point is equal to the agent's income in the current period, to the case of a general reference point which is unique to the individual. The key mechanism is that there is a kink in utility at the reference point which creates a discontinuity in the marginal utility of consumption, so that two individuals who are otherwise similar will exhibit different savings behaviour due to differences in their reference points. The model is first developed as a 2 period model with no uncertainty, which is amenable to closed-form solution for the savings function. Appendix A.1 extends some key intuitions of the model to cases with income uncertainty and to T periods.

The individual's utility function follows Koszegi and Rabin (2006, 2007, 2009) and is given by:

$$U(c_{it}, r_{it}) = u(c_{it}) + \mu(u(c_{it}) - u(r_{it})) \quad (1.1)$$

where:

- c_{it} = consumption of individual i in period t
- r_{it} = reference point of individual i in period t

for each individual in the sample $i = 1, \dots, n$ and each period $t = 1, \dots, T$. The function (1.1) contains a standard utility component over consumption for that period $u(c_{it})$, where $u'(c_{it}) > 0$ and $u''(c_{it}) < 0$ for all c_{it} i.e. the utility function is concave. There is also an additively separable reference-dependent component with an individual-specific reference point r_{it} , which is defined as $\mu(u(c_{it}) - u(r_{it}))$, where $\mu(\cdot)$ is a linear loss aversion function such that:

$$\mu(x) = \begin{cases} \gamma x, & \text{if } x \geq 0 \\ \lambda x, & \text{if } x < 0 \end{cases} \quad (1.2)$$

with $\lambda > \gamma$. Thus, the individual compares the utility they receive from consumption to the utility they would receive if they were to consume at their reference point. Due to the kink in equation (1.2) they experience loss aversion: the loss in utility λ from

consuming an amount which is below their reference point by Δx is higher than the gain in utility γ from consuming an amount which is above their reference point by Δx .

The individual solves a simple consumption-savings problem over 2 periods, with utility function (1.1) in each period. Dropping the i subscript for simplicity, the maximisation problem is given by:

$$\begin{aligned} \max_{c_1, c_2} \quad & U(c_1, r_1, c_2, r_2) = U(c_1, r_1) + U(c_2, r_2) \\ \text{subject to} \quad & c_1 + c_2 = y_1 + y_2 \end{aligned} \tag{1.3}$$

where:


- y_t = income in period t
- saving $s_1 \equiv y_1 - c_1$

The individual receives income y_t in each period and is free to save or dissave between periods, with savings s_1 representing consumption deferred to period 2. It is assumed that we can abstract from the interest rate r and discount rate ρ in the main analysis, or equivalently that $\frac{1+r}{1+\rho} = 1$. However, relaxing this assumption does not affect the main predictions of the model as there is still a discontinuity in marginal utility for any given interest rate and discount factor.

Figure 1.7 shows the individual's indifference curves, which are kinked owing to reference-dependence and loss aversion. The slope of the indifference curves vary depending on whether the individual is above or below their reference point in each period. In each period t , when $c_t = r_t$ there is a kink in the indifference curve, marking the transition across distinct marginal rates of substitution between periods. In both panels, the upper segment of the indifference curve is for the case where $c_2 \geq r_2$, but $c_1 < r_1$. The middle segment is when both $c_2 < r_2$ and $c_1 < r_1$, while the lower segment is where $c_2 < r_2$, but $c_1 \geq r_1$. Since the budget constraint binds and the reference points are assumed to be sufficiently high in these figures, it is not possible for both $c_1 > r_1$ and $c_2 > r_2$.

As a result of these kinks the reference point in each period acts as an attractor, pulling consumption towards it. As the marginal utility of consumption is discontinuous at the reference point, individuals can be 'stuck' at a reference point with no incentive to move either way. Furthermore, *ceteris paribus* a change in r_1 relative to y_1 can induce completely different consumption smoothing behaviour. This is shown

Figure 1.7: Indifference Curves with Reference Points



Indifference curves for a consumption-smoothing individual with a utility function featuring reference-dependent preferences and loss aversion. The budget line is given by B and its slope is $-(1+r)$. The individual receives income y_t and has a reference point r_t in each period $t = 1, 2$ and at each reference point there is a kink, owing to a discontinuity in the marginal utility of consumption. In the top panel the first period reference point is r_1^a and the individual maximises utility at (c_1^a, c_2^a) , receiving utility U_1 . Since $y_1 < c_1$, they dissave in period 1 and save in period 2. In the bottom panel the first period reference point is $r_1^b < r_1^a$ individual maximises utility at (c_1^b, c_2^b) , receiving utility $U_2 > U_1$. Since $y_1 > c_1$ and they save in period 1 and dissave in period 2.

in the difference between the top and bottom panels in Figure 1.7: in the top panel, the reference point is $r_1^a > y_1$ and the individual maximises utility at U_1 , dissaving. In the bottom panel, all else is equal but the reference point is $r_1^b < y_1$. As a result the new location of the kink at the reference point shifts the marginal utility of consumption in period 1 downwards and the individual saves, maximising utility at $U_2 > U_1$.

The following proof clarifies this prediction and shows it can be tested in two different ways. Substituting $c_2 = y_1 + y_2 - c_1$ into the objective function (1.3) yields:

$$U(c_1) = u(c_1) + \mu(u(c_1) - u(r_1)) + u(y_1 + y_2 - c_1) + \mu(u(y_1 + y_2 - c_1) - u(r_2)) \quad (1.4)$$

The derivative of $\mu(\cdot)$ is discontinuous. With a slight abuse of notation, write the derivative of $\mu(f(x))$ with respect to x as:

$$\mu'_t(f'(x)) = \begin{cases} \gamma f'(x_t), & \text{if } f(x_t) \geq 0 \\ \lambda f'(x_t), & \text{if } f(x_t) < 0 \end{cases} \quad (1.5)$$

Assume that $u(c_t) = \ln(c_t)$. The first order condition with respect to c_1 can be expressed as:

$$\begin{aligned} U'(c_1) &= \frac{1}{c_1} + \frac{\mu'_1}{c_1} - \frac{1}{y_1 + y_2 - c_1} - \frac{\mu'_2}{y_1 + y_2 - c_1} = 0 \\ \iff \frac{1 + \mu'_1}{c_1} &= \frac{1 + \mu'_2}{y_1 + y_2 - c_1} \\ \iff c_1 &= \frac{(1 + \mu'_1)(y_1 + y_2)}{2 + \mu'_1 + \mu'_2} \end{aligned} \quad (1.6)$$

Since savings $s_1 \equiv y_1 - c_1$, it can be shown that:

$$s_1 = \frac{(1 + \mu'_2)y_1 - (1 + \mu'_1)y_2}{2 + \mu'_1 + \mu'_2} \quad (1.7)$$


Result (1.7) can be explained intuitively: savings are increasing period 1 income and decreasing in period 2 income. If the individual is below their reference point in period 1, this reduces the proportion of period 1 income saved *and* the penalty of period 2 income, as they consume more now to avoid losses in the present. If the individual is below their reference point in period 2, this increases the proportion of period 1 income saved *and* the penalty of period 2 income, as they save more now to avoid losses in

the future. The predicted savings function is shown in Figure 1.8, and in line with the indifference curves in 1.7 there is a point at which consumption is ‘stuck’ at the reference point, as well as a discontinuity after this point. A change in the reference point is represented by a shift in the grey dotted line, which causes the individual with first period income \tilde{y}_1 to jump from one part of the savings function to another.

Let us summarise the implications of this model as a whole in intuitive terms before proceeding to empirical estimation. The presence of reference-dependent preferences and loss aversion creates a kink in the utility function and a discontinuity in the marginal utility of consumption. This means shifts in r_1 can lead to jumps in savings behaviour as depicted in Figures 1.7 and 1.8. According to the model, individuals who are above their reference points will be more likely to be on the upper section of their savings function, therefore saving more, while those who are below their reference points will be on the lower section of their savings function, therefore saving less. If we assume a particular functional form for utility as in (1.7), we can characterise the savings function precisely: the coefficient on y_1 is higher when $c_1 > r_1$, so in certain scenarios the individual will even switch from borrowing to saving if r_1 is higher. Appendix A.1 extends the model to the cases of income uncertainty and to T periods. Though it is not possible to generate a closed form solution for a savings function in these cases, the prediction of a discontinuity in the savings function, as well as a region in which individuals are stuck at this discontinuity, can be generalised.

This has two related implications. Firstly, on average savings levels for individuals who are above their reference point should be higher than savings levels for individuals who are below their reference point – what is termed the ‘adaptation effect’ in this chapter, when the reference point depends on income in the previous period. Secondly, those who are above their reference point should save less out of each unit of income, holding their future income constant – a test of whether the adaptation effect is driven by loss aversion. Using the change SFS variable as a measure of where individuals are relative to their reference point, we can translate this prediction into empirical estimations.

Figure 1.8: Savings Function with Loss Aversion



Shows the savings function from the theoretical result in (1.7) for a given r_2 and y_2 , with first period reference point equal to r_1^a in the top panel and r_1^b in the bottom panel, with $r_1^a < r_1^b$. Values of savings for which $c_1 = r_1$ are shown by the grey dotted line and where this intersects the lower portion of the savings function it causes the function to flatten. This is because consumption is 'stuck' at r_1 : moving upward to the new savings function would not be feasible as the discrete jump in savings would push consumption below the reference point. Where the upper portion of the savings function intersects the line $c_1 = r_1$, income is sufficient for savings to jump upwards while retaining $c_1 \geq r_1$ and savings is now given by the upper portion of the function. An individual with first period income \tilde{y}_1 is marked by a black dot in each panel. As the reference point increases there is a shift outward in the dotted line which creates a jump in savings behaviour between panels. Due to this difference between r_1^a and r_1^b the individual saves at s_1^a and s_1^b respectively, a shift in savings behaviour similar to the indifference curve diagram in Figure 1.7.

1.5 Methodology

This section will detail the estimation techniques used to test the predictions of the model. Self-reported savings are regressed on the SFS variables, which are coded as dummy variables with a dummy for each answer category — using the middle answers as the base dummies³. The category of individuals who perceive a negative change can be considered below their reference point (represented by $\mathbb{1}(c_{it} < r_{it})$):

$$sav_{it} = \beta_1 \mathbb{1}(c_{it} < r_{it}) + \beta_2 y_{it} + \beta_3 \mathbb{1}(c_{it} < r_{it}) \cdot y_{it} + \delta X_{it} + h(t) + \bar{u}_i + u_{it} \quad (1.8)$$

where i indexes individuals and t indexes years, X_{it} is a set of covariates including the current and expectations SFS variables, and $h(t)$ represents a full set of year dummies. The rich dataset means that a wide range of covariates can be controlled for including income, age, job status, region and home ownership status. The error term is split into a standard idiosyncratic error u_{it} and an individual fixed effect \bar{u}_i , which controls for individual characteristics which are correlated with permanent differences in the level of both savings and the SFS variable. For example, some individuals may have an especially negative outlook on their finances, causing them to save diligently but also to evaluate their financial situation negatively, and this would bias the estimation.

As illustrated by equation (1.7) in Section 1.4, loss aversion can be detected by interacting income with the change SFS variable so that there are separate coefficients for those at or above, versus those below their reference points – assuming change SFS measures consumption relative to the reference point. The base MPS is given by β_2 while the coefficient β_3 gives the deviation in the MPS for those below their reference points, and $H_0 : \beta_3 = 0$ against the one-sided alternative $H_A : \beta_3 < 0$ tests for the existence of loss aversion.

The change SFS variable is assumed to give a window into $r_1 = r(y_0, w_0)$, where y_0 is the individual's income last period, w_0 is other components of wealth, and $r(\cdot)$ is a function which captures the individual's subjective judgment. This is a key assumption: although r_1 is compared to the level of consumption in the model, it is not a function of the previous period's consumption. In addition, it must be assumed here that the expectations SFS variable controls for c_2 relative to r_2 .

The idea that reference points may depend on income instead of, or as well as consumption is not uncommon in the literature. One's pay cheque (Cohen, 2015) or past

³'Just about getting by' for Current Financial Situation, 'About the same' for Expected Financial Situation, and 'About the same' for Change in Financial Situation.

pay cheque Dellavigna et al. (2017) serving as a reference point is plausible as this number is salient to an individual. This chapter goes one step further and argues that other components of wealth might affect reference points directly, for instance someone who outright owns a large house is likely to have higher aspirations and therefore a higher reference point than someone who rents a small flat. Theoretically an exogenous rather than endogenous variable is useful as a determinant of reference points because it prevents the starvation mechanic discussed earlier, where the individual postpones consumption to reduce their reference point in the future — a problem found in Bowman et al. (1999) and discussed in detail in Michaelides (2002). Nevertheless, it is possible that somebody whose financial situation has improved (declined) could still be consuming below (above) their reference point, so this assumption is a limitation of the present analysis.

There are two estimators used for regression (1.8). The first is a Tobit model which treats those who report saving zero as censored dissaving and makes assumptions about the distribution of this dissaving in the estimation (Tobin, 1958). This is consistent with the theoretical model outlined above and utilises the most information. The second, alternative estimation technique is to use standard OLS which treats the reported zeroes for savings as genuine instances of no saving at all. This infers that those who report zero saving are not actually saving or dissaving, which is appropriate if analysis is restricted to ‘active’ saving: as discussed above this may be a conscious, qualitatively different decision to not saving. The main results are not sensitive to these choices about model specification.

If censoring is allowed then the standard Tobit estimator is inconsistent in the presence of fixed effects (FE), making Honoré (1992)’s Trimmed Least Absolute Deviations (LAD) estimation more suitable. Unlike the standard Tobit estimator Trimmed LAD does not require parametric assumptions on the error terms so the non-normality of savings in Figure 1.2 is not a problem. It is also robust to heteroskedasticity across individuals. Both Tobit FE and standard FE regression will be reported throughout for robustness. Note that although there is a high degree of censoring there is also a high sample size and Honoré (1992)’s own simulations show that the estimator is consistent in high sample sizes even with a large degree of censoring.

Since income is included in the covariates it is reasonable to ask how to interpret the coefficient on income in regressions (1.8) when controlling for current SFS, which should itself take income into account. There are two interpretations. Firstly, two individuals may have different incomes but perceive their financial situation differently,

as per the relative income hypothesis (Duesenberry, 1949). A rich individual may have a higher standard for what constitutes ‘living comfortably’ than a poor individual, but their differing incomes will mean the former saves more than the latter. Secondly, two individuals may evaluate their financial situations similarly although they have different compositions. One individual might be ‘living comfortably’ because they are retired in a large house with no mortgage, but relatively low income, while another might be ‘living comfortably’ because they are an investment banker with both high income and high outgoings. Yet the former will probably be dissaving while the latter will be saving, and this will be picked up by the differences in income between the two.

One limitation of the SFS variables is reverse causality: it could be argued that an individual is now above their reference point because they saved – or below it because they did not – contrary to regression (1.8). As documented earlier the SFS variable which accounts for the reason why SFS changed, including because of changes in savings behaviour, can be included in the regressions to control for the latter possibility.

1.6 Results

This section reports the results from the regressions used to estimate to test whether savings function are characterised by loss aversion, as well as some additional empirical analysis to investigate the determinants of reference points. Table 1.7 shows results from regression (1.8), for standard fixed effects regressions in columns (1) and (2) and Honoré (1992)’s Tobit fixed effects estimator in columns (3) and (4). Every SFS variable is reported for completeness and as expected, having a good current financial status is associated with more savings, while a bad current financial status is associated with less savings once censoring is taken into account. Negative expectations are associated with higher savings, in line with “saving for a rainy day”, though positive expectations are associated with less savings only in standard fixed effects, suggesting bias when censoring is not taken into account.

The main coefficients of interest are those for the decline in SFS, which is taken as an indication that an individual is below their reference point, as well as the interaction of the decline in SFS with income. As above, a loss averse utility function implies that those below r will have lower savings levels for a given level of income, which is driven by their lower marginal propensity to save. Therefore the coefficient on the decline in SFS should be negative, as should the interaction term. Additionally, decline

SFS should not be statistically distinct from zero once the interaction term is included in the regression.

In Table 1.7 columns (1) and (3) show the coefficient on the decline in SFS is negative for both estimators without the interaction term. In columns (2) and (4) the interaction term is also negative and statistically distinct from 0 at all conventional levels of significance. The hypothesis of loss aversion is therefore borne out by the data. For the standard fixed effects specification in column (2) the inclusion of the interaction term reverses the sign of the coefficient on SFS decline, which is not a feature of the model above. Despite this, in column (4) the same coefficient is insignificant when censoring is taken into account, further supporting the predictions of the model from (1.7) and indicating that the Tobit fixed effects model is the appropriate specification.

The effects reported in Table 1.7 are economically substantial: in column (3), those below their reference point save £750 per year less, the equivalent of having almost £7,000 more income. In column (4) the marginal propensity to save for individuals below r is 40% lower than for those at or above r . The coefficient on income is itself comparable macroeconomic estimates on savings rates. An increase in income of £1,000 increases annual savings by £112, about 11%. According to the ONS, the average savings rate across the sample period was 10.2%, which is surprisingly similar to this result. Appendix A.2 repeats the main regression with lagged values of the change SFS variable to further account for endogeneity concerns, and though the coefficients are smaller the results are consistent with the main analysis.

To get an idea of the macroeconomic magnitude of these effects, according to a back-of-the envelope calculation if everyone in the sample who is currently below their reference point were instead at or above it, the savings rate would increase by two percentage points (to 13%). This is more than a 20% increase and could have important implications for policy: for example, if individuals who are below their reference point are more likely to spend then fiscal stimulus may be more effective during recessions, especially if it is targeted toward groups who have been most affected. It would also mean recessions have larger welfare effects than previously supposed, a possibility also explored by De Neve et al. (2018) using subjective well-being data in Europe.

1.6.1 Unfulfilled Expectations

Reference points could depend on unfulfilled expectations rather than simple changes in financial situation. For example, if an individual expected an improvement last year but experienced roughly no change then they might consider themselves below their

Table 1.7: Effect of Reference Points on Savings

	FE		FE Tobit	
	(1)	(2)	(3)	(4)
Good Financial Status	226.0*** (10.87)	228.6*** (10.71)	979.0*** (58.16)	996.0*** (57.54)
Bad Financial Status	38.10*** (9.920)	3.990 (9.270)	-729.0*** (118.2)	-791.1*** (117.4)
Positive Expectations	-40.85*** (13.84)	-39.75*** (13.82)	-60.85 (46.06)	-58.75 (45.56)
Negative Expectations	128.3*** (16.33)	119.5*** (16.38)	329.5*** (59.58)	311.8*** (58.36)
SFS Declined	-215.0*** (10.76)	111.1** (43.53)	-748.8*** (48.59)	130.3 (135.4)
Income (1,000s)	28.95*** (4.381)	34.05*** (5.036)	112.2*** (12.55)	120.5*** (12.44)
Income * SFS Declined		-24.00*** (3.328)		-48.98*** (8.368)
Observations	196443	196443	196443	196443

a. This table shows regressions of self-reported savings on the subjective financial situation (SFS) variables and covariates. Columns (1) and (2) use a standard fixed effects specification while columns (3) and (4) use Honoré (1992)'s fixed effects Tobit estimator to compensate for censoring. Whether an individual is below r is indicated by whether they report their Subjective Financial Situation (SFS) as having gotten worse, which is interacted with income to test whether the savings function is characterised by loss aversion. The base category are those who report their situation has improved or stayed the same. Covariates include income, year dummies, and demographic variables.

b. Standard errors in parentheses, clustered by individual.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

reference point and the change SFS variable might not pick this up. Lagged expectations are measured by the expected SFS variable at time $t - 1$ and each answer can be interacted with answers from the change SFS variable at time t to show whether an individual's expectations have been fulfilled or not. Table 1.8 shows these interaction dummies. There are 9 possibilities in total, which can be categorised into 3 different cases: falling short of expectations, exceeding expectations, and matching expectations.

Table 1.8: SFS interactions summary statistics

Type of change	Proportion in Sample
Panel A: Matches Expectations	
Expected Improvement → Realised Improvement	12.6 %
Expected Decline → Realised Decline	5.5 %
Expected No Change → Realised No Change	38.7 %
Total	51.5%
Panel B: Exceeds Expectations	
Expected Decline → Realised Improvement	1.5 %
Expected Decline → Realised No Change	3.5 %
Expected No Change → Realised Improvement	11.9 %
Total	15.5%
Panel C: Falls Short of Expectations	
Expected Improvement → Realised Decline	5.1 %
Expected Improvement → Realised No Change	9.0 %
Expected No Change → Realised Decline	12.2 %
Total	24.4%

Shows the percentages in the sample for whether individual's past expectations — as measured by their expected SFS in the previous year — are realised, as measured by their change SFS in the present year. Panel A shows those who matched their expectations; Panel B shows those who have exceeded their expectations; Panel C shows those who have fallen short of their expectations. There are nine possibilities in total, with three in each panel. Values are shown only for those whose expectations SFS the previous year and change SFS this year are both available ($n=185,448$). Percentages may not add up to exactly 100 due to rounding.

Table 1.9 shows the results from regressions (1.8), only with the aggregate dummies for falling short of expectations included (exceeding or matching expectations is the base category). The results are analogous to those in Table 1.7: in columns (1) and (3) falling short of expectations has a negative effect on savings, while in (2) and (4) the interaction of falling short of expectations and income also has a negative effect. The inclusion of the interaction term reverses the sign of the coefficient on falling short of expectations for standard fixed effects, while the same coefficient becomes insignificant

Table 1.9: Effect of Unfulfilled Expectations on Savings

	FE		FE Tobit	
	(1)	(2)	(3)	(4)
Income (1,000s)	28.74*** (4.490)	32.93*** (5.107)	108.5*** (12.59)	114.8*** (12.61)
Unexpected Negative Change	-182.0*** (17.00)	90.03** (42.44)	-658.7*** (62.87)	-32.78 (124.9)
Unexpected Negative Change * Income		-18.87*** (3.046)		-34.70*** (7.216)
Observations	170465	170465	170465	170465

a. This table shows regressions of self-reported savings on the subjective financial situation (SFS) variables and covariates. Columns (1) and (2) use a standard fixed effects specification while columns (3) and (4) use Honoré (1992)'s fixed effects Tobit estimator to compensate for censoring. Whether an individual is below r is indicated by whether whether their reported change in subjective financial situation falls short of their expected change in financial situation last year, which is also interacted with income to test whether the savings function is characterised by loss aversion. Covariates include the original SFS variables, year dummies, and demographic variables.

b. Standard errors in parentheses, clustered by individual.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

for the Tobit fixed effects estimator. Overall these results suggest that unfulfilled expectations are also a candidate for being below one's reference point, and the effects are of similar magnitude to the previous section.

1.6.2 Factor Analysis

The previous analysis used the SFS variables to measure reference points, imposing that the change SFS variable measured reference points in the context of financial changes. In this section a more flexible approach is taken: it is assumed that all three SFS variable could capture some aspects of reference points but the data is allowed to determine the relative importance of each variable. This is done by using factor analysis, which parses down multiple observed variables in same category (in this case, subjective financial situation) to a smaller set of orthogonal variables which measure the latent underlying processes that determine those variables.

It is found that the three SFS variables can be reduced to two processes, one of which is interpreted as subjective reporting style and the other of which is interpreted as a separate negative component of financial situation. It is shown that the former has a positive impact on savings, illustrating that subjective perceptions affect consumption-smoothing. Conversely, the second factor has a negative impact on savings, consistent with losses having a distinct negative impact from gains as in prospect theory (Kahneman and Tversky, 1979) and also the psychology literature (Argyle, 1987; Rajabi and Sheykhabani, 2009; Bayliss et al., 2014). Further analysis on couples who live together illustrates that although partners differ in their reported SFS, partner's SFS as measured by the two factors does not have predictive power for the household head's savings.

The three SFS variables are used in a polychoric factor analysis, which is designed for discrete variables (with missing values dropped, leaving a total sample size of 220,000). The results are reported directly from Stata in Table 1.10 and because two factors have eigenvalues greater than one, the third factor can be omitted. Together the two factors explain over 82% of the variation in the SFS variables, with Factor 1 accounting for almost half and factor 2 accounting for a further third of the total variation.

Table 1.11 reports the factor loadings on each SFS variable – to what extent the factors correlate with the original variables – which enables interpretation of the factors in terms of the original variables. Factor 1 unambiguously seems to be measure of favourable subjective financial perceptions, since it is positively correlated with all SFS variables. Indeed, since it is correlated with positive expectations as well as positive current financial situation, it is best interpreted as a measure of optimism or subjective reporting style. Further support for this interpretation is that when the 12-item mental health & well-being GHQ questionnaire is included in the factor analysis (unreported),

Table 1.10: Factor Analysis of SFS

Factor	Eigenvalue	Proportion Explained	Cumulative Proportion Explained
Factor 1	1.465	0.488	0.488
Factor 2	1.014	0.338	0.827
Factor 3	0.520	0.173	1.000

Polychoric factor analysis of the three subjective financial situation (SFS) variables. The second column reports the eigenvalues from the linear transformation of the variables, which reflect the explanatory power of each factor. If the eigenvalue is greater than one then the factor has sufficient predictive power. The third columns shows that proportion of variation in the original variables explained by each factor, while the fourth column shows the same, but cumulatively.

it returns the same factors and better mental well-being is positively correlated with Factor 1.

Factor 2 is harder to interpret as it is negatively correlated with current SFS, positively correlated with expectations, but uncorrelated with changes. Table 1.12 reports the cross-tab of the original SFS variables once more to help understand these correlations and it shows negative expectations are correlated with positive current financial situation, suggesting a regression to the mean phenomenon where individuals who are currently ‘finding things difficult’ expect them to improve again. This explains the negative correlation with current SFS and positive correlation with expectations. Column 3 of Table 1.12 also shows that positive versus negative changes are roughly equally likely for individuals experiencing good and bad current financial situations, which explains why Factor 2 (a measure of a bad financial situation) would be net uncorrelated with change.

Table 1.11: Factor Correlations with Original Variables

Factor	Current SFS	Future SFS	Change SFS
Factor 1	0.4954	0.2360	0.5976
Factor 2	-0.1838	0.2880	0.0386

Correlations of the two factors with the original SFS variables. Factor 1 is positively correlated with all three variables while Factor 2 is negatively correlated with current SFS, positively correlated with future SFS and uncorrelated with change SFS.

Factor 2 is therefore best interpreted as a separate negative component of financial situation, consistent with it being negatively correlated with current SFS. Further support for Factor 2 as a negative measure is found in that it’s positively correlated with unemployment and mental health strain (GHQ) but negatively correlated with having

Table 1.12: Cross-Correlations of SFS Variables

Current SFS	Future SFS	Change SFS
Good	2.011	2.211
Medium	2.282	2.346
Bad	2.394	2.192

Correlations between the three SFS variables with the sample restricted to cases where none are missing (219,629 observations). Change SFS is positively correlated with both current and expected SFS, but current SFS is uncorrelated with expected SFS.

a degree and being married. Such an interpretation is consistent with reference dependent preferences and loss aversion as the negative component is qualitatively different to the positive component, a foundation for loss aversion which has been proposed in the psychology literature (Diener, 1984; Argyle, 1987).

Table 1.13: PCA Analysis

	FE			Tobit FE		
	(1)	(2)	(3)	(4)	(5)	(6)
Real Income (Logs)	87.69*** (6.759)	78.29*** (6.618)	76.56*** (6.632)	1652.5*** (182.0)	1565.1*** (184.5)	1522.6*** (182.9)
Factor 1		160.7*** (6.674)	152.7*** (7.250)		1861.3*** (133.6)	1835.0*** (135.3)
Factor 2			-67.27*** (14.22)			-1037.6*** (195.5)
Observations	216798	216798	216798	216798	216798	216798

a. This table shows regressions of self-reported savings on the factors extracted from the subjective financial situation (SFS) variables using polychoric factor analysis. Columns (1), (2) and (3) use a standard fixed effects specification while columns (4), (5) and (6) use Honoré (1992)'s fixed effects Tobit estimator to compensate for censoring. Covariates include income, year dummies, and demographic variables.

b. Standard errors in parentheses, clustered by individual.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.13 shows the results of regressing self-reported savings on the two factors and covariates. In all specifications Factor 1 is strongly significant and has a positive effect on savings: in the most complete specification a one standard deviation increase in Factor 1 increases savings by £969. Similarly, Factor 2 has a strong negative effect on savings: in the most complete specification, a one standard deviation increase in Factor 2 decreases savings by £266. These effects are all conditional on income, which retains a strongly significant and economically substantial positive effect on savings

in all specifications where it is included. The regression results therefore imply that subjective perceptions of financial situation have a positive effect on realised savings, while ‘negative financial situation’ exerts an independent influence on savings, which is consistent with a gain-loss utility function as in (1.3).

Within-Household SFS

In couple households budget constraints are often joint and therefore savings is not only an individual decision. In the context of reference points it is logical to ask whether one’s partner’s reference point affects one’s own savings. For example, if the household head’s financial situation is steady but their partner has experienced a negative financial shock there is a possibility the head will change their behaviour to account for this, and vice-versa. Formally we can say the household reference point $R_{it} = R(r_{it}, r_{jt})$ where i represents an individual and j represents that individual’s partner. In this section it is shown that partners report different levels of SFS around 30% of the time but, despite this, partner’s SFS – as measured by the variables extracted from factor analysis – does not impact one’s own savings.

The sample is restricted only to the heads of households of 2 who are married or cohabitating, around 30,000 person-years. Table 1.14 shows around 11,000 report different SFS to their partners – 30% of the time – within these households. Panel A shows all households in the restricted sample, while Panel B shows these disagreements persist in similar proportions even if the sample is restricted to those households which report that resources are shared equally. This further supports the notion that it is subjective perceptions rather than intra-household distribution which is driving this heterogeneity. This is an important distinction between the relative and permanent income hypotheses: in the former, subjective perceptions matter for savings while in the latter they do not. That the SFS variables both depend on subjective perceptions and affect savings behaviour therefore supports the relative over the permanent income hypothesis.

Table 1.15 shows the same regression from Table 1.13, but restricted only to the heads of households and including the partner’s SFS. Columns (1) and (3) replicate the results from the previous table in the smaller sample, at the cost of higher standard errors, while columns (2) and (4) add partner’s SFS. Results unambiguously show that partner’s SFS has no predictive power of the head’s savings: Factor 1 has very large standard errors, while Factor 2 is not even identified. While there may be concerns that this is due to lower statistical power, note that the coefficients for the original

Table 1.14: Disagreements in SFS Within Households

Household Head		Partner		
		Better	About the Same	Worse
	Panel A: All Partner Households			
	Better	4,336 (13.8%)	1,960 (6.2%)	721 (2.3%)
	About the Same	1,925 (6.1%)	12,974 (41.3%)	1,937 (6.2%)
	Worse	648 (2.1%)	2,089 (6.7%)	4,792 (15.3%)
	Panel B: Report Sharing Resources Equally			
	Better	669 (16.1%)	265 (6.4%)	100 (2.4%)
	About the Same	252 (6.1%)	1,481 (35.64%)	270 (6.5%)
	Worse	107 (2.57%)	301 (7.2%)	711 (17.1%)

Disagreements in SFS between partners in couple households. Panel A shows all partner households (N=31,382) while Panel B shows only those households who answer the question “How are your household finances organised?” with “We share resources equally”. Within each panel, each row contains a different household head’s answer to the change SFS question while each column contains their partner’s. Therefore the diagonals within each panel represents times partners’ and household head’s SFS answers are the same while all cells outside the diagonal represent times they are different. The latter occurs around 30% of the time in total.

factors are still statistically significant, and that the descriptives in Table 1.14 show there is substantial variation in own versus partner’s SFS. Thus results suggest that partner’s subjective perceptions do not affect one’s own savings. Appendix A.2 runs a robustness check by using the partner’s change SFS instead of the extracted factors, and it is consistent with the results presented here.

Table 1.15: PCA Analysis

	FE		FE Tobit	
	(1)	(2)	(3)	(4)
Factor 1 (Own)	39.37** (18.81)	39.75* (21.11)	576.1** (249.3)	542.2** (228.9)
Factor 2 (Own)	-93.48** (46.40)	-93.52** (46.42)	-1003.6** (489.7)	-994.1** (472.8)
Real Income (Logs)	22.26 (16.91)	22.27 (16.89)	564.7 (580.4)	561.4 (581.4)
Factor 1 (Partner's)		-0.707 (20.39)		9.863 (188.1)
Factor 2 (Partner's)		0 (.)		0 (.)
Observations	30970	30970	30970	30970

a. This table shows regressions of self-reported savings on the factors extracted from the subjective financial situation (SFS) variables using polychoric factor analysis. Columns (1) and (2) use a standard fixed effects specification while columns (3) and (4) use Honoré (1992)'s fixed effects Tobit estimator to compensate for censoring. All regressions include individual fixed effects, year dummies and demographic covariates.

b. Standard errors in parentheses, clustered by individual.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.7 Conclusion

The results in this chapter establish that past financial situation acts as a reference point, inducing a lower marginal propensity to save due to loss aversion. Subjective Financial Situation (SFS) variables were proposed as an empirical measure of reference points, allowing individuals' subjective perceptions to determine their own reference points. Savings functions were investigated in the British Household Panel Survey (BHPS), which contains data on both individuals savings and the SFS variables. Regressions showed that individuals below their reference point save substantially less – equal to about £7,000 less income – than those above their reference point, and that this is driven by a 40% lower marginal propensity to save. These results are consistent with a model of consumption and savings featuring reference-dependent preferences and loss aversion. Additionally, it was shown that using unfulfilled expectations as a measure of reference points yields similar results.

Further analysis decomposed the SFS variables into two latent factors, one of which was interpreted as reporting style and the other as the negative component of SFS. The former illustrates that the SFS variables capture subjective perceptions, while the latter is consistent with loss aversion (Kahneman and Tversky, 1979), and the psychology literature's interpretation of losses as psychologically distinct from gains (Diener, 1984; Argyle, 1987). Regressions of savings on the two factors show that the first has a positive impact while the latter has a negative impact on savings, as would be predicted by these interpretations. Additionally, it was shown that within-household differences in perceptions about SFS between partners are quite common, supporting the idea that the reference points are partly subjective. Nevertheless, regressions restricted to the heads of couple of households but including the partner's factors showed that only one's own SFS affects savings.

This chapter contributes to both the life cycle literature and the prospect theory literature, which still lacks go-to measures of reference points in applied settings. The findings in this chapter should be interpreted as evidence that research on income changes, and in consumption and savings models more generally, should work to incorporate the effect of loss aversion on shifts in income — as well as the 'objective' effects the literature typically focuses on such as transitory versus permanent, aggregate versus idiosyncratic, and expected versus unexpected shocks (Jappelli and Pistaferri, 2010; Attanasio and Weber, 2010). Moreover, the combination of reference-dependent preferences and loss aversion creates an important asymmetry in savings behaviour between individuals who are above versus below their reference points which may be

important for understanding policy analysis of fiscal multipliers, or the welfare costs of recessions (De Neve et al., 2018). This effect is macroeconomically consequential: if everyone in the sample who is currently below their reference point were instead at or above it, the savings rate would increase by two percentage points.

Another issue raised by this chapter is that Subjective Financial Situation (SFS) variables can be a valuable tool for investigating how individual's perceptions affect their behaviour. Theoretical models such as the one in Section 1.4 can help to disentangle what the SFS variables measure when compared to objective variables like income, providing a way to test reference-dependence in household datasets which has not been fully utilised in the past. Versions of the SFS variables are available in several datasets in Germany, Australia, the US and the Netherlands so there is ample room for future research to utilise them more extensively.

Appendices

A.1 Proof with Uncertainty

This section extends the model in Section 1.4 to the case of uncertainty, first for 2 periods with a simple type of uncertainty and then for T periods with a more general type of uncertainty. It is shown that the prediction of the adaptation effect can be generalised, though it is not possible to generate a closed-form savings function as in equation (1.7).

2 Periods, With Uncertainty The introduction of uncertainty to the 2-period model does not change the logic of the proof. In fact, the introduction of a range of possible states eliminates the reference point condition for the second period because expected values ‘smooth’ over the kink. To illustrate this, suppose second period income can have two states:

$$y_2 = \begin{cases} y_H \equiv y_1 + \varepsilon, & \text{with probability } p \\ y_L \equiv y_1 - \varepsilon, & \text{with probability } 1 - p \end{cases} \quad (9)$$


where $y_H \geq r_2 > y_L$. This means that in the ‘good’ state of the world the individual receives sufficient income to consume at or above their reference point without having saved; in the ‘bad’ state of the world the individual receives insufficient income to consume at or above their reference point without having saved.

In order to understand savings behaviour we should evaluate the marginal utility of consumption at y_1 . If the marginal utility of consuming at y_1 is lower than the marginal utility of consuming at $E[y_2]$, the individual will save, and vice-versa for dissaving. Consider the condition for whether to save in period 1 with second period income uncertain:

$$\begin{aligned} u'(y_1)(1 + \mu') &< pu'(y_H)(1 + \gamma) + (1 - p)u'(y_L)(1 + \lambda) \\ \iff u'(y_1) &< pu'(y_H)\frac{1 + \gamma}{1 + \mu'} + (1 - p)u'(y_L)\frac{1 + \lambda}{1 + \mu'} \\ \iff y_1 &> u'^{-1}\left(pu'(y_H)\frac{1 + \gamma}{1 + \mu'} + (1 - p)u'(y_L)\frac{1 + \lambda}{1 + \mu'}\right) \end{aligned} \quad (10)$$

This will depend on whether y_1 is above, at, or below r_1 , which determines the value of μ' . However, there is no longer a need for a condition specifying y_2 relative to r_2 because the convex combination of outcomes eliminates it. Thus, perhaps unexpectedly, the addition of uncertainty makes this prediction of the model clearer. This is

Figure 9: Indifference Curves with Reference Points with Uncertainty in Period 2



Indifference curve for a consumption-smoothing individual with a utility function featuring reference-dependent preferences and loss aversion. The budget line is given by B and its slope is $-(1+r)$. The individual receives income y_t and has a reference point r_t in each period $t = 1, 2$, but in period 2 income y_2 is stochastic. This ‘smooths’ over the kink for period 2, since some realisations will be above and others below the reference point. Therefore there is a kink only at r_1 , owing to a discontinuity in the marginal utility of consumption. The individual maximises utility at (c_1, c_2) , receiving utility U_3 .

illustrated in Figure 9, where there is only one kink in the marginal rate of substitution between periods 1 and 2. As in the main text, this kink produces the adaptation effect as a result of loss aversion, so savings levels will be higher for individuals above their reference points, owing to a higher marginal propensity to save.

On the other hand, the addition of uncertainty means that a closed form solution for the savings function itself is no longer possible. This is a well known result in the life cycle literature: even in absence of reference-dependent preferences, CRRA or log utility makes solving for optimal consumption-savings impossible and the models have to be approximated numerically (Browning and Lusardi, 1996).

T Periods, With Uncertainty. This section generalises the model to T periods with a general process for stochastic income, showing that it yields the same prediction as the simpler versions of the model due to the kink in marginal utility created by the combination of reference-dependent preferences and loss aversion. Our individual now maximises:

$$\max_{c_t} \sum_{t=1}^T \left(\left(\frac{1}{1+\rho} \right)^t u(c_t) + \mu(u(c_t) - (r_t)) \right) \quad (11)$$

$$s.t. \quad a_{t+1} = (1+r)(y_t + a_t - c_t)$$

Where a_t are assets in period t , r is the real interest rate, c_t is consumption in period t , and y_t is stochastic income which is distributed i.i.d. $y_t \sim N(m, \sigma_y^2)$ with mean m and standard deviation σ_y .

Proposition 1A. *Consider the individual's decision to save in period 1. For some future period $t = \tau$, there exist subjective savings thresholds $\phi(y_\tau, r_\tau)$ and $\psi(y_\tau, r_\tau)$, where both $\phi(y_\tau, r_\tau)$ and $\psi(y_\tau, r_\tau)$ are increasing in y_τ and with $\phi(y_\tau, r_\tau) < \psi(y_\tau, r_\tau)$, such that:*

If $y_1 > r_1$, then the individual will save if $y_1 > \phi(y_\tau, r_\tau)$, dissave if $y_1 < \phi(y_\tau, r_\tau)$ and neither save nor dissave if $y_1 = \phi(y_\tau, r_\tau)$

If $y_1 < r_1$, then the individual will save if $y_1 > \psi(y_\tau, r_\tau)$, dissave if $y_1 < \psi(y_\tau, r_\tau)$ and neither save nor dissave if $y_1 = \psi(y_\tau, r_\tau)$

If $y_1 = r_1$, then the individual will save if $y_1 > \psi(y_\tau, r_\tau)$, dissave if $y_1 < \phi(y_\tau, r_\tau)$ and neither save nor dissave if $\psi(y_\tau, r_\tau) > y_1 > \phi(y_\tau, r_\tau)$.

Proof: Consider the agent's decision of whether to save or dissave in period 1. The marginal utility at y_1 is:

$$MU_1 = u'(y_1)(1 + \mu') \quad (12)$$

In order for the agent to save in period 1, there must be at least one period τ where the agent borrows. The expected marginal utility of income in period τ is:

$$\begin{aligned} MU_\tau &= \int_{-\infty}^{+\infty} u'(y_\tau)(1 + \mu') dy_\tau \\ &= \int_{-\infty}^{r_\tau} u'(y_\tau)(1 + \lambda) dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau)(1 + \gamma) dy_\tau \end{aligned} \quad (13)$$

For an individual to save, it must be the case that $MU_1 < MU_\tau$:

$$u'(y_1)(1 + \mu') < \int_{-\infty}^{r_\tau} u'(y_\tau)(1 + \lambda) dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau)(1 + \gamma) dy_\tau \quad (14)$$

Analogous to above, assume that $y_1 > r_1$:

$$\begin{aligned} u'(y_1)(1 + \gamma) &< \int_{-\infty}^{r_\tau} u'(y_\tau)(1 + \lambda) dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau)(1 + \gamma) dy_\tau \\ \iff u'(y_1) &< \int_{-\infty}^{r_\tau} u'(y_\tau) \frac{(1 + \lambda)}{(1 + \gamma)} dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau) dy_\tau \\ \iff y_1 &> u'^{-1} \left(\int_{-\infty}^{r_\tau} u'(y_\tau) \frac{(1 + \lambda)}{(1 + \gamma)} dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau) dy_\tau \right) \equiv \phi(y_\tau, r_\tau) \end{aligned} \quad (15)$$

As in the proof for the 2 period deterministic case, we can repeat this exercise when $y_1 < r_1$. In that case the analogue of (15) is that the individual will save if $y_1 > u'^{-1} \left(\int_{-\infty}^{r_\tau} u'(y_\tau) dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau) \frac{(1 + \gamma)}{(1 + \lambda)} dy_\tau \right) \equiv \psi(y_\tau, r_\tau)$. When $y_1 = r_1$ the individual will once again be caught on the 'kink'. This leads directly to Corollary 1A.

Corollary 1A. Suppose $\psi(y_\tau, r_\tau) > y_1 > \phi(y_\tau, r_\tau)$, with $y_1 > r_1$. The individual will save a positive amount, as this will not incur losses in period 1 but will reduce or prevent losses in period τ .

Now suppose, ceteris paribus, that r_1 is higher, so that $y_1 < r_1$. The individual will now dissave in order to smooth out the losses over both periods.

Suppose instead that $y_1 = r_1$. The individual will neither save nor dissave.

Proof: Following the above, if $y_1 > r_1$, the individual will save iff:

$$y_1 > u'^{-1} \left(\int_{-\infty}^{r_\tau} u'(y_\tau) \frac{(1+\lambda)}{(1+\gamma)} dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau) dy_\tau \right) \equiv \phi(y_\tau, r_\tau) \quad (16)$$

Since by assumption $y_1 > \phi(y_\tau, r_\tau)$, the individual saves.

Now consider a rise in r_1 , *ceteris paribus*, such that $y_1 < r_1$. The individual will dissave iff:

$$y_1 < u'^{-1} \left(\int_{-\infty}^{r_\tau} u'(y_\tau) dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau) \frac{(1+\gamma)}{(1+\lambda)} dy_\tau \right) \equiv \psi(y_\tau, r_\tau). \quad (17)$$

Since by assumption $y_1 < \psi(y_\tau, r_\tau)$, the individual dissaves.

Finally, consider the case where $r_1 = y_1$. The individual will save iff:

$$y_1 > u'^{-1} \left(\int_{-\infty}^{r_\tau} u'(y_\tau) dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau) \frac{(1+\gamma)}{(1+\lambda)} dy_\tau \right) \equiv \psi(y_\tau, r_\tau). \quad (18)$$

and dissave iff:

$$y_1 < u'^{-1} \left(\int_{-\infty}^{r_\tau} u'(y_\tau) \frac{(1+\lambda)}{(1+\gamma)} dy_\tau + \int_{r_\tau}^{+\infty} u'(y_\tau) dy_\tau \right) \equiv \phi(y_\tau, r_\tau) \quad (19)$$

Since by assumption $\psi(y_\tau, r_\tau) > y_1 > \phi(y_\tau, r_\tau)$, the individual consumes all of y_1 , neither saving nor dissaving.

Corollary 1A shows that individuals who have the same income and expectations will demonstrate different savings behaviour depending on whether their first period income is at, above or below their first period reference point. Thus, the differences in r_1 between individuals create a ‘regime shift’ that causes a sudden change in behaviour. This point does not apply so starkly to all individuals: those for whom $\psi(y_\tau, r_\tau) < y_1$ or $y_1 < \phi(y_\tau, r_\tau)$ will not experience this change. Intuitively, this is because if the individual’s income increase (decrease) between periods 1 and 2 is sufficiently large, they will dissave (save) no matter their reference point. However, this ‘kink’ in savings behaviour can be demonstrated for any initial value of c_1 and c_2 .


A.2 Robustness Checks

This section presents two additional sets of results which serve as a check on the main analysis by supplementing regression (1.8): firstly, it uses lagged values of the

change SFS variables to account for endogeneity problems; secondly, it uses the partner's raw SFS variables instead of the factors extracted from polychoric factor analysis.

Lagged SFS. Because the savings variable and the change SFS variable refer to the same period there may be contemporaneous circumstances and behaviours which contribute to both SFS and savings. Although the main analysis used a variable which asks why financial situation changed to control for reverse causality (see Table 1.6), it is possible that these answers were not inclusive of all possibilities. Another approach is to lag change SFS by one period so that the individual's position relative to their reference point precedes their savings behaviour, in line with the model in Section 1.4. The logic of this choice is shown in Figure A1.

Figure A1: Timeline of savings behaviour



The timeline of saving when using the lagged, as opposed to contemporaneous SFS variable. This series of events is consistent with the model in Section 1.4 and less prone to concerns about endogeneity.

Table A1 shows results from regression (1.8) but using the change SFS variable lagged by one period – all other covariates, including the other SFS variables, are contemporaneous. Results are qualitatively consistent with the main analysis, showing that a decline in SFS reduce savings and the interaction between the decline and income is also negative. The reason this specification is not used in the main analysis is because the lagged report of SFS may not be salient to individuals in the present year. Since this chapter is concerned with individual's perceptions, this is an important drawback. Additionally, as the results are consistent with one another reverse causality is clearly not an overriding problem.

Partner's Change in SFS. The main analysis tested whether the factors extracted from polychoric factor analysis on the three SFS variables had an effect on savings. It then restricted the sample to partner households and tested whether or not one's partner's factors had an effect on one's own savings. It is natural to ask whether one's partner's reference point itself has a direct effect on one's own savings by using the raw change SFS variable rather than the extracted factors. Regression model (20) shows the regression with partner's SFS variables included, where i indexes the household

Table A1: Effect of Reference Points on Savings

	Savings	
	(1)	(2)
SFS Declined (lagged)	-107.1*** (12.19)	117.7*** (43.90)
Income (1,000s)	28.46*** (4.495)	30.91*** (4.879)
SFS declined (lagged)* Income		-17.18*** (3.401)
Observations	170465	170465
Controls		

a. Regressions of self-reported savings on the lagged subjective financial situation (SFS) variables and covariates. The base category are those who reported their situation had stayed the same. Covariates include year dummies, the other SFS variables for the present period, and demographic variables.

b. Standard errors in parentheses, clustered by individual.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

head and j is their partner:

$$\begin{aligned}
 sav_{it} = & \beta_1 \mathbb{1}(c_{it} < r_{it}) + \beta_2 y_{it} + \beta_3 \mathbb{1}(c_{it} < r_{it}) \cdot y_{it} + \\
 & + \zeta_1 \mathbb{1}(c_{jt} > r_{jt}) + \zeta_2 \mathbb{1}(c_{jt} > r_{jt}) \cdot y_{it} + \delta X_{it} + h(t) + \bar{u}_i + u_{it}
 \end{aligned} \tag{20}$$

Table A2 shows the results from regression (20), estimated just on the household heads of 2 adult households (married or cohabiting couples) with no other household members. Columns (1) and (2) recreate columns (1) and (2) from Table 1.7 for the smaller sample and the results are qualitatively unchanged. A decline in one's own SFS reduces savings, while income increases savings, and the interaction of income and a decline in SFS is negative. Columns (3) and (4) add the concurrent variables for one's partner and a decline in partner's SFS has no effect on savings (the coefficients cannot even be identified). This is consistent with the factor regressions in Table 1.15.

Table A2: Effect of Reference Points on Savings

	Own SFS		Partner's SFS	
	(1)	(2)	(3)	(4)
Own SFS Declined	-214.7*** (50.96)	368.4*** (106.5)	-214.7*** (50.96)	368.4*** (106.5)
Income (1,000s)	28.99*** (9.741)	48.06*** (9.516)	28.99*** (9.741)	48.06*** (9.516)
Own SFS Declined* Income		-41.75*** (8.416)		-41.75*** (8.416)
Partner's SFS Declined			0 (.)	0 (.)
Partner's SFS Declined* Income				0 (.)
Observations	10572	10572	10572	10572

a. Regressions of self-reported savings on both the individual's and their partner's subjective financial situation (SFS). Both columns use Honoré (1992)'s fixed effects Tobit estimator to compensate for censoring. Whether an individual is above or below r is indicated by whether they report their Subjective Financial Situation (SFS) as having gotten better or worse, respectively. The base category are those who report their situation has stayed the same. All regressions include covariates such as income, year dummies, and demographic variables.

b. Standard errors in parentheses, clustered by individual.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 2

Making Up for Losses: Estimation of Marginal Utility with Reference-Dependent Preferences

1 Introduction

Reference-dependent preferences have become a key tool in economists' arsenal and have been used in a variety of applications. They contrast with classical utility functions because they introduce relative comparisons, whereas classical utility is based on absolute consumption or wealth positions. When combined with loss aversion, one implication of reference-dependent preferences is that the marginal utility of consumption will be higher for individuals who are below their reference point than for those who are above it. This paper contributes to the literature by testing this assumption for the first time using 'experienced utility' in the form of subjective well-being (SWB) and food consumption as a proxy for total consumption.

A standard model featuring reference-dependent preferences and loss aversion is tested using British panel data. Parameter estimates are consistent with the null hypothesis of loss aversion and indicate that it is characterised by a smooth utility function rather than a discontinuous one. The analysis favours a reference point based on habit formation (Pollak, 1970) as opposed to relative income (Duesenberry, 1949). However, the estimated parameters suggest model misspecification. Additional analysis decomposes the SWB measure into positive and negative affect to test the hypothesis that loss aversion is largely driven by negative affect as in psychological theories where losses are qualitatively distinct from gains (Argyle, 1987), but this hypothesis is not

supported. Instead, it appears that food consumption in general has a larger impact on negative components of well-being.

Economists increasingly use self-reported SWB in order to measure experienced utility (Kahneman et al., 1997), an area known as the economics of happiness literature (Frey and Stutzer, 2002, 2010, 2018). This complements the more traditional, ordinal approach to utility by allowing researchers to ask direct questions about what affects agents' utility. Whereas ordinal utility asks how preferences are revealed through actions, cardinal approaches such as SWB ask agents to state their preferences through their own subjective evaluation of their situation. Over the past few decades there has been an emerging consensus over the usefulness of SWB. Early literature concerned what exactly SWB measured (Kimball and Willis, 2006); whether it was robust to repeated testing of the same individuals (Diener, 1984; Argyle, 1987); and how it correlated with standard survey variables such as employment, income, age, marriage and so forth (Frey and Stutzer, 2002, 2010). Now that researchers have convinced themselves that SWB gives sensible and reliable results, it has been used to test more novel hypotheses such as the influence of religion, war, and most importantly for this paper, behavioural heuristics and biases (Frey and Stutzer, 2018).

A key feature of economic models is that agents derive utility from consumption, equating the marginal utility of consumption to the marginal cost of a commodity, or to the marginal disutility of labour. Traditional models assume that marginal utility falls with income, an assumption for which there is extensive evidence (Horowitz et al., 2007). Models with reference-dependent preferences complement this by assuming the marginal utility of consumption varies depending on whether agents are above or below their reference point. If an agent is below their reference point they experience 'loss aversion', which simultaneously acts as a penalty for their total utility and augments their marginal utility of consumption (Kahneman and Tversky, 1979). These are the key implications of loss aversion investigated in this paper and will be discussed in more detail later.

There are a number of potential determinants of reference points proposed in the literature (Barberis, 2013). Models of habit formation build up a stock of bonus utility from past consumption, creating an adaptation lag as individuals adjust to changes (Pollak, 1970; Constantinides, 1990). Bowman et al. (1999) build a reference-dependent model with $c_{i,t-1}$ as an argument of the reference point, which creates a similar effect.

Adaptation has been documented in the BHPS in a number of contexts including adaptation to subjective health status (Groot, 2000), wage levels (Clark, 1999), and incurring a disability (Oswald and Powdthavee, 2008). Guariglia and Rossi (2002) favour a habit-formation model over a standard model using consumption data in the first few waves of the BHPS. The previous chapter of this thesis showed that adaptation was present in short-term consumption-smoothing behaviour using the BHPS dataset and a subjective measure of reference points.

Alternatively, reference points can be thought of as a social comparison, an idea originally theorised by Duesenberry (1949). In this case $r_{it} = \hat{c}_t$, representing consumption averaged over some relevant peer group such as colleagues, peers, or region. Several papers have proposed that SWB – experienced utility – depends on relative comparisons, beginning with Easterlin (1974). He famously observed that although rich individuals are happier than poor individuals within a given country, countries do not experience nationwide increases in happiness as GDP grows (past the threshold of GDP necessary for basic needs). This is hypothesised to be because people compare their consumption to those around them, which grows with their own, neutralising utility gains. Although debate continues over the so-called ‘Easterlin Paradox’ (Stevenson and Wolfers, 2008; Easterlin et al., 2010), the idea that utility depends on comparisons to those around you has become more widely accepted among SWB researchers (Clark and Oswald, 1995; Clark et al., 2008; Frey and Stutzer, 2010).

In this paper both habit formation and relative comparisons will be considered as potential determinants of reference points so that similarly to Bowman et al. (1999) $r_{it} = \phi \hat{c}_t + (1 - \phi)c_{i,t-1}$, with $0 \leq \phi \leq 1$. This paper is the first in the literature to test the assumption that marginal experienced utility depends on consumption relative to a reference point, and this also allows an investigation of the value of ϕ i.e. which, if either, reference point best explains SWB. Thus the paper contributes to both the literature on the economics of happiness and the literature on the determinants of reference points in field settings.

This paper is conceptually close to Finkelstein et al. (2012), who estimate the impact of health status on marginal utility as measured by SWB. They demonstrate a large decline in marginal utility for individuals who experience negative health status, which is consequential for models predicting the demand for healthcare and lifecycle savings. They note that “even a moderate amount of state dependence can have a substantial effect on the conclusions of such calculations. Moreover, not only the magnitude but also the sign of any potential state dependence is a priori ambiguous.” Similar logic

applies here: the theory of reference-dependent preferences assumes a positive impact of loss aversion on marginal utility, but neither its existence nor magnitude has been demonstrated using experienced utility, both of which have consequences for policy and welfare analysis. Again, a simple example was the first chapter of this thesis, where reference-dependent preferences had a macroeconomically consequential impact on aggregate savings.

The chapter proceeds as follows. Section 2 outlines the dataset used, the harmonised British Household Panel Survey/Understanding Society (BHPS/USoc). Section 3 details the theoretical specification and how it translates into a regression model. Section 4 presents and discusses the regression results. Section 5 concludes, including implications for future work.

2 Data

The data cover the period 1991-2016 in the United Kingdom (UK) and are a combination of the British Household Panel Survey (BHPS) – detailed in the previous chapter – and subsequent Understanding Society (USoc) survey, which continued to interview many of the same people (Institute for Social and Economic Research et al., 2019). These surveys provide detailed information on social, economic and demographic variables, allowing for the investigation of the influence of a wide range of variables on well-being. The data are longitudinal and the total sample size is around 610,000 person-years, with 105,000 unique persons and an average time in the survey of around 6 years.

The USoc survey is funded by the ESRC and multiple government departments from across the UK and is run in combination with the University of Warwick, the University of Essex, and the London School of Economics. USoc is designed as a follow up to the BHPS and thus far covers the period 2009-2016, having begun the year after the BHPS was discontinued. It contains similar data to the BHPS, retaining many key variables but is much larger, interviewing 40,000 households each year while retaining much of the original BHPS sample. Like the BHPS, only adults are given the main survey, and the data are gathered using a similar combination of methods (Knies, 2014). The result is that the two surveys are highly comparable in many respects.

One issue with using both datasets in the same analysis is that the BHPS cohort did not join the USoc survey until 2010, which means that in the combined dataset everyone from the BHPS disappears in 2009, and then some of them reappear starting

from 2010. This interrupts the long and continuous panel. However, there is still a large cross section available for 2009, and many of the people surveyed for the first time in 2009 continue into subsequent years.

General Health Questionnaire (GHQ) The GHQ is considered a valid instrument for detecting psychological distress (Argyle, 1987). However, the 12-item version has increasingly been used as an indicator of general well-being in the literature (Clark and Oswald, 1995; Bayliss et al., 2014; Jones and Wildman, 2008; Oswald, 1997). Higher GHQ scores have been associated with a number of behaviours that increase among unemployed individuals such as watching TV casually, sitting at home, smoking and window shopping, while lower GHQ scores have been associated with behaviours that decrease amongst unemployed individuals (Warr, 1984). The GHQ-12 questionnaire used in the BHPS has also been shown to be robust to long term retest effects, showing it is a reliable psychological instrument (Pevalin, 2000).

Both the BHPS and USoc contain the GHQ-12 questionnaire. This contains 12 questions which attempt to elicit an individual's mental well-being, which are as follows:

1. Have you recently been able to concentrate on whatever you're doing?
2. Have you recently lost much sleep over worry?
3. Have you recently felt that you were playing a useful part in things?
4. Have you recently felt capable of making decisions about things?
5. Have you recently felt constantly under strain?
6. Have you recently felt you couldn't overcome your difficulties?
7. Have you recently been able to enjoy your normal day to day activities?
8. Have you recently been able to face up to problems?
9. Have you recently been feeling unhappy or depressed?
10. Have you recently been losing confidence in yourself?
11. Have you recently been thinking of yourself as a worthless person?
12. Have you recently been feeling reasonably happy, all things considered?

A moment of inspection shows that there are two types of question: one is positively framed while the other is negatively framed. All 12 questions are answered on a 4-point scale, with a higher score indicating that the individual is in a worse mental state.

Table 2.1: Examples of Question Types from GHQ Questionnaire

	Positive Question Example	Negative Question Example
Value	“Have you recently lost much sleep over worry?”	“Have you recently been feeling reasonably happy, all things considered?”
1	Not at all	More than usual
2	No more than usual	Same as usual
3	Rather more	Less so
4	Much more	Much less

Two different questions from the GHQ-12 questionnaire, one of which is positively phrased while the other is negatively phrased. Although the questions have different possible answers, in each case a higher value indicates increased mental distress. Every question in the GHQ-12 fits into one of these two categories.


This is illustrated in Table 2.1, which uses the question about sleep and the question about overall happiness as examples of positive and negative questions, respectively. I will, however, multiply all measures by -1 so that a higher score indicates improved SWB and vice-versa, which makes the regression coefficients easier to interpret later on. Figure 2.1 shows the overall distribution of modified GHQ (Likert) answers in the sample, and most answers are in the upper half of the distribution.

The GHQ questions can be taken individually or combined, and the BHPS data already contain two commonly used ways to summarise all 12 questions. I will primarily use the Likert scale, which is a 0 to 36 point scale that simply sums the responses to each question coded 0 (no mental health problem) to 3 (serious mental health problem), so that a score of 36 indicates high mental distress and a score of 0 indicates no mental distress. Item (12), shown in Table 2.1, is also a more straightforward measure of ‘happiness’ so will be used as a robustness check.

Another measure of life satisfaction is also available in the dataset, which asks “How satisfied are you with your life overall, on a scale of 1-7?”. It is not used in favour of the GHQ-12 for several reasons. Firstly, it is not available for waves 1-5 and 11 of the BHPS while the GHQ-12 has a long and consistent panel of 25 years. Secondly, it is single-item as opposed to multi-item and the latter has been shown to be more reliable to retest validity (Argyle, 1987). Thirdly, as this paper looks at the marginal utility of consumption, the short-term nature of the GHQ-12 questions is actually more appropriate — and is in line with the food question, which is also over a short-term time frame.

Figure 2.2 shows both GHQ measures over the sample period, each split into four categories, and all are steady prior to the financial crisis except for the proportion who are happy in the GHQ (Likert) measure, which declines slightly. The main shift is

Figure 2.1: GHQ (Likert) Distribution




The distribution of GHQ (Likert) in the BHPS/USoc sample, multiplied by -1 so low values indicate low well-being and high values indicate high well-being.

between 2008-2009 where the proportion of people who are unhappy actually falls in both graphs, while the proportion who are happy falls a little. Both ‘happy’ and ‘unhappy’, the two middle categories, rise subsequently. Previous literature has documented that SWB is quite stable in the UK over the Great Recession (Bayliss et al., 2014), though not that it increases. Importantly, Understanding Society replaced the BHPS in 2009 so this increase may just be an artefact of the change in the sample and survey. This can be accounted for by year dummies in the regression analysis.

Food Consumption. Food consumption is reported throughout the sample and asks “Thinking about your weekly food bills approximately how much does your household usually spend in total on food and groceries?”. It is banded in every wave except the first, and bands are shown in Table 2.2. Fintel (2006) shows that such banding does not affect estimation and inference, particularly when the bands are tight as in the present case. General descriptives for all variables including food consumption and GHQ are shown in Table 2.3.

Figure 2.2: Subjective Well-Being Over Sample Period



The behaviour of both the GHQ (Likert) measure (top) over the sample period, from 1991-2016. The Likert measure is split into four categories manually, where less than -20 is coded as 'Very Unhappy', between -20 and -10 is coded as 'Unhappy', between -10 and -5 is coded as 'Happy', and over -5 is coded as 'Very Happy'.

Table 2.2: Food Expenditure Bands for Households in BHPS/USoc

Band (£)	%
0	3.35
Under 10	0.46
10-19	3.04
20-29	7.57
30-39	10.70
40-49	12.21
50-59	14.43
60-79	19.51
80-99	12.69
100-119	8.86
120-139	3.79
140-159	2.02
160+	1.378

Household's estimated weekly expenditure on food in the combined BHPS/USoc dataset from 1992-2016 for 228,836 households in total. Answers are given in bands, and the first wave is excluded as it was the only wave where households reported raw amount.

Table 2.3: Economic and demographic variables in the combined BHPS/USoc dataset

	Mean	Standard Deviation
Panel A: Standard Variables		
Food consumption	1737.37	950.47
GHQ (Likert)	11.15	5.45
Age	46.59	18.64
Number of persons in household	2.92	1.48
Number of own children in household	0.50	0.93
Real income (£)	15,562	16,275
Hours worked (full sample)	16.48	18.44
Hours worked (employed)	33.5	12.1
Subjective health status (1-5 scale)	2.4	1.07
Panel B: Category Variables		
Male	0.46	0.50
White	0.78	0.41
Household		
Own/Mortgage house	0.70	0.46
Rent house	0.25	0.43
Marital Status		
Living as couple	0.35	0.48
Married	0.35	0.48
Widowed	0.03	0.18
Divorced	0.03	0.17
Qualifications		
Degree or higher	0.18	0.39
Some qualifications	0.61	0.49
No qualifications	0.17	0.38
Employment Status		
Self-employed	0.07	0.26
Employed	0.48	0.50
Unemployed	0.05	0.21
Retired	0.22	0.41
Student	0.07	0.25
Long term sick/disabled	0.04	0.19
Region		
London & Southeast	0.21	0.41
North	0.24	0.43
Midlands & East	0.32	0.46
Wales & Southwest	0.12	0.33
Scotland & Channel Islands	0.10	0.30
Northern Ireland	0.01	0.10

The means and standard deviations of variables in the combined BHPS/USoc dataset from 1991-2016, measuring economic and demographic characteristics for the entire sample ($N = 612,771$). Panel A variables are all positive numbers, while Panel B variables are measured in proportions between 0 and 1. All given moments exclude missing values for that variable.

3 Theoretical Framework

Consider a standard utility function featuring a reference point and loss aversion which follows the previous chapter in using the specification of Koszegi and Rabin (2006, 2007, 2009):

$$U(c_{it}, r_{it}) = u(c_{it}) + \mu(u(c_{it}) - u(r_{it})) \quad (2.1)$$

Where $i = 1, \dots, N$ indexes individuals and $t = 1, \dots, T$ indexes years. Consumption is given by c_{it} with a reference point r_{it} . Once more $\mu(\cdot)$ is a linear loss aversion function such that:

$$\mu(x) = \begin{cases} \gamma x, & \text{if } x \geq 0 \\ \lambda x, & \text{if } x < 0 \end{cases} \quad (2.2)$$

with $\lambda > \gamma$ i.e. loss aversion, which drives the results tested in this chapter. The individual compares the utility they receive from consumption to the utility they would receive if they were to consume at their reference point, and experiences loss aversion: due to the parameter λ , the loss in utility from consuming an amount which is below their reference point by Δx is higher than the gain in utility from consuming an amount which is above their reference point by Δx .

If we assume we can proxy for utility U with subjective well-being measures, it is straightforward to translate this theoretical model into a regression model. Dropping the i 's and t 's for convenience, let $u(c) = \alpha \cdot \log(c)$. It follows that:

$$U(c, r) = \alpha \cdot \begin{cases} \log(c) + \gamma(\log(c) - \log(r)), & \text{if } \log(c) \geq \log(r) \\ \log(c) + \lambda(\log(c) - \log(r)), & \text{if } \log(c) < \log(r) \end{cases} \quad (2.3)$$


Let $\bar{c} \equiv \log(c)$ and $\bar{r} \equiv \log(r)$, and add an error term $\varepsilon \sim N(0, \sigma^2)$:

$$U = \begin{cases} \alpha\bar{c} + \alpha\gamma(\bar{c} - \bar{r}) + \varepsilon, & \text{if } \bar{c} \geq \bar{r} \\ \alpha\bar{c} + \alpha\lambda(\bar{c} - \bar{r}) + \varepsilon, & \text{else} \end{cases} \quad (2.4)$$

Define the indicator $d = \mathbb{1}(\bar{x} < \bar{r})$. In a regression framework the model is now:

$$U = d \cdot \alpha[\bar{c} + \gamma(\bar{c} - \bar{r})] + (1 - d)\alpha[\bar{c} + \lambda(\bar{c} - \bar{r})] + \varepsilon \quad (2.5)$$

Figure 2.3: Loss Aversion with a Kink versus with a Jump



Two different log-linear utility functions featuring reference-dependent preferences and loss aversion. The function U_1 features a kink at the reference point but is continuous. To the left of $\log(r)$ the slope is $(1 + \lambda)$ and to the right it is $(1 + \gamma)$. Function U_2 is identical but for a discontinuity at $\log(r)$ as utility jumps upwards. For exposition purposes it is assumed that $\alpha = 1$ in this graph.

We can rearrange (2.5) to collect terms:

$$U = \alpha \bar{r} + \alpha(1 + \gamma)(\bar{c} - \bar{r}) + \alpha(\lambda - \gamma) \cdot d \cdot (\bar{c} - \bar{r}) + \varepsilon \quad (2.6)$$

There are two testable hypotheses implied by (2.6). The first is the hypothesis of no loss aversion $H_0 : \gamma = \lambda$ against the one-sided alternative of loss aversion $H_A : \gamma < \lambda$. This can be tested by whether the coefficient $\alpha(\lambda - \gamma)$ attached to $d \cdot (\bar{c} - \bar{r})$ is significantly different to zero.

The second hypothesis is less obvious and is illustrated in Figure 2.3 which shows two different utility functions: U_1 , which has only a change in slope; and U_2 , which also has a discontinuity. This is the distinction between a continuous loss aversion function as in (2.2) and one with a discrete ‘notch’ in utility as discussed in Blackburn and Chivers (2015); Seibold (2017); Genicot and Ray (2017). It is worth investigating which, if either, is the most empirically accurate assumption. We can test whether the utility function is truly ‘kinked’ – continuous, with a change in slope at $\log(r)$ – or whether there is a discontinuity as well as a kink by adding an additional variable d to (2.6) and whether or not the coefficient is statistically significant¹. If it is significant, the equation corresponds to the dotted line in Figure 2.3.

As mentioned earlier, this chapter tests two potential determinants of reference

¹Note that it is theoretically possible there is a discontinuity but no kink, which is not shown in Figure 2.3.

points, habit formation and relative income:

$$r_{it} = \phi \hat{c}_t + (1 - \phi) c_{i,t-1} \quad (2.7)$$

The regression specifications used in this chapter are equivalent to inserting (2.7) into (2.5) for the values $\phi = \{0, 1\}$. Put simply, it is just a case of using two different measures of the reference point in separate regressions.

4 Results

The theoretical model above is translated into regression specification (2.8) by using SWB as a proxy for utility, food consumption as a proxy for total consumption, and both the habit formation and relative income measures of reference points. Individual fixed effects ω_i are added to control for individual heterogeneity, along with demographic covariates (see Table 2.3), and year dummies $h(t)$. The final regression specification is:

$$\begin{aligned} SWB_{it} = & \omega_i + \alpha \ln(r_{it}) + \alpha(1 + \gamma)(\ln(c_{it} - \ln(r_{it}))) + \beta d + \\ & + \alpha(\lambda - \gamma) \cdot d \cdot (\ln(c_{it} - \ln(r_{it}))) + \delta X_{it} + h(t) + \epsilon_{it} \end{aligned} \quad (2.8)$$

where i indexes individuals and t indexes time, while c_{it} is the consumption of individual i and r_{it} is their reference point. The indicator function $d \equiv \mathbb{1}(c_{it} < r_{it})$ is equal to one if the individual is below their reference point and 0 otherwise. The regression is run separately for the two different measures of r_{it} : consumption in the previous period $c_{i,t-1}$, and average regional consumption \hat{c}_t . Unfortunately it is not possible to include both reference point measures in the same regression due to low variation in the consumption measure.

Results are shown in Table 2.4, with two separate panels for the habit formation and relative income specifications. Column (1) strictly estimates the model parameters only, while column (2) includes the dummy d to test for the discontinuity. Note that in both panels this dummy is insignificant, suggesting that the discontinuity is not present in this case, and its inclusion does not qualitatively affect the other estimated coefficients. The null hypothesis of no loss aversion is $H_0 : \lambda - \gamma = 0$, which can be tested by whether the estimate $\hat{\alpha}(\hat{\lambda} - \hat{\gamma})$ is statistically distinct from 0. In Panel A this coefficient is positive and significant so the null can be rejected, with $p < 0.05$ against the one sided alternative $H_A : \lambda < \gamma$ in both columns.

Table 2.4: Effect of Reference Points on Marginal Utility of Consumption

	GHQ (Likert)	
	(1)	(2)
Panel A: Habit Formation Reference Point		
$\hat{\alpha}$	0.103** (0.0490)	0.105** (0.0491)
$\hat{\alpha}(1 + \hat{\gamma})$	-0.0128 (0.0549)	0.0344 (0.0649)
$\hat{\alpha}(\hat{\lambda} - \hat{\gamma})$	0.191** (0.0793)	0.168** (0.0812)
$\hat{\beta}$		0.0435 (0.0282)
Observations	454389	454389
Panel B: Relative Income Reference Point		
$\hat{\alpha}$	0.133*** (0.0381)	0.129*** (0.0383)
$\hat{\alpha}(1 + \hat{\gamma})$	0.132*** (0.0381)	0.129*** (0.0383)
$\hat{\alpha}(\hat{\lambda} - \hat{\gamma})$	-0.00002 (0.000018)	-0.0002 (0.00018)
$\hat{\beta}$		-0.309 (0.308)
Observations	588097	588097

a. This table shows regressions of the GHQ (Likert) subjective well-being index on measures of food consumption in the combined BHPS/USoc dataset. The specification is as in (2.8), with the estimated coefficients returning estimates of the theoretical parameters α, λ, β , and γ . Panel A uses a reference point of last year's consumption – habit formation – while panel B uses average regional consumption — relative income. All regressions include individual fixed effects, year dummies and demographic covariates.

b. Standard errors in parentheses, clustered at the individual level

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Therefore in Panel A the evidence points to loss aversion. However, there is a complication: although the estimate of α is positive and significant as would be expected, the estimate of $\alpha(1 + \gamma)$ is small in magnitude and insignificant. This would seem to imply that γ is negative, which is contradictory to the theory and common sense as it would suggest that people have ‘gain aversion’. It is therefore likely that the model is misspecified, or possibly that there is not enough statistical power in our chosen measure of consumption. Overall, the panel provides suggestive evidence that the data support the idea of habits acting as a reference point, but a more nuanced theoretical framework may be needed.

In Panel B, when the reference point is defined as relative income, the empirical findings are inconsistent with the presence of reference-dependent preferences and loss aversion. The good news is that the estimates are more consistent with one another in indicating a standard concave utility function. Once again $\hat{\alpha}$ is positive and actually of a similar magnitude to Panel A. The coefficient $\hat{\alpha}(\hat{\lambda} - \hat{\gamma})$ is tiny and statistically insignificant at all conventional levels, so there is no loss aversion and $\lambda = \gamma$. Furthermore, the coefficient $\hat{\alpha}(1 + \hat{\gamma})$ is positive and significant but is statistically indistinguishable from $\hat{\alpha}$, which suggests that $\gamma = 0$. Taken together, these empirical findings are consistent with a model where consumption has a positive effect on utility, but there is no gain-loss utility or reference-dependence for relative income reference points.

4.1 Decomposing GHQ

It has been argued in the psychology literature that “good is not merely the opposite of bad”: bad experiences are evaluated on a different metric to good experiences, such that the absence of bad does not imply good and vice versa (Argyle, 1987). Another, related argument is that psychological tests focus too much on the negative, and thus that measures such as GHQ are more likely to pick up bad feelings than good ones (Diener, 1984). In a working paper, Bayliss et al. (2014) follow Rajabi and Sheykhshabani (2009) by splitting the GHQ-12 into two equally sized components, one of which reflects ‘positive affect’ and one of which reflects ‘negative affect’. The questions used for each GHQ measure are shown in Table 2.5.

Table 2.6 shows regression (2.8) separately for the positive and negative GHQ components. The results are qualitatively consistent across the positive and negative components of GHQ, but the magnitudes are all substantially higher for the negative component. While it is tempting to interpret this as evidence of loss aversion affecting

Table 2.5: Bayliss et al. (2014) Affect Decomposition

GHQ Positive Affect Questions	GHQ Negative Affect Questions
Have you recently...	Have you recently...
1) ...been able to concentrate on whatever you're doing?	1) ...lost much sleep over worry?
2) ...felt that you were playing a useful part in things?	2) ...felt constantly under strain?
3) ...felt capable of making decisions about things?	3) ...felt you couldn't overcome your difficulties?
4) ...been able to enjoy normal day-to-day activities?	4) ...been feeling unhappy or depressed?
5) ...been able to face up to problems?	5) ...been losing confidence in yourself?
6) ...been feeling reasonably happy, all things considered?	6) ...been thinking of yourself as a worthless person?

Decomposition of the GHQ-12 questionnaire into questions which represent positive affect and questions which represent negative affect, in line with Bayliss et al. (2014). Each category contains half of the total number of questions.

negative GHQ more, a closer look shows that this is not the case. *All* coefficients in columns (3) & (4) are larger in magnitude than columns (1) & (2), including $\hat{\alpha}$, which is simply the coefficient governing $u(c_{it})$. Therefore the correct interpretation is that that in the sample, individual's negative utility is more strongly affected by (food) consumption than positive utility, which also scales up the estimate $\hat{\alpha}(\hat{\lambda} - \hat{\gamma})$. This makes sense because food is crucial for survival, and those lacking food are likely to experience the negative well-being effects in the GHQ such as lack of sleep, concentration, and unhappiness.

In fact, if anything loss aversion is more pronounced for the positive component of GHQ. If we take the point estimates from column (1) – bearing in mind $\hat{\alpha}$ is highly uncertain – we find that $(\hat{\lambda} - \hat{\gamma}) \approx 2$, while in column (3) $(\hat{\lambda} - \hat{\gamma}) \approx 1$. Another unexpected finding is that in column (4), the coefficient $\hat{\beta}$ is positive and significant, implying that there is a discontinuity at the reference point. Thus, for negative GHQ the utility function is better characterised by U_2 than U_1 in Figure 2.3.

Table 2.6: Effect of Reference Points on Marginal Utility of Consumption

	GHQ (Positive)		GHQ (Negative)	
	(1)	(2)	(3)	(4)
$\hat{\alpha}$	0.0216* (0.0124)	0.0217* (0.0124)	0.131*** (0.0229)	0.132*** (0.0229)
$\hat{\alpha}(1 + \hat{\gamma})$	0.0132 (0.0137)	0.0150 (0.0160)	0.0300 (0.0246)	0.0575** (0.0287)
$\hat{\alpha}(\hat{\lambda} - \hat{\gamma})$	0.0457** (0.0196)	0.0448** (0.0199)	0.137*** (0.0351)	0.123*** (0.0358)
$\hat{\beta}$		0.00163 (0.00717)		0.0249** (0.0126)
Observations	412707	412707	412712	412712

a. This table shows GHQ decomposed into positive and negative components as in Bayliss et al. (2014), with each component separately regressed on measures of food consumption in the combined BHPS/USoc dataset. The specification is as in (2.8), with the estimated coefficients returning estimates of the theoretical parameters α , λ , β , and γ . All regressions use a habit formation (lagged consumption) measure for reference points and include individual fixed effects, year dummies, and demographic covariates.

b. Standard errors in parentheses, clustered at the individual level

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5 Conclusion

This paper has proposed a simple test of a theoretical model with reference-dependent preferences and loss aversion. Following the economics of happiness literature by using data on subjective well-being as a measure of ‘experienced utility’ (Kahneman et al., 1997), as well as food consumption, the model was tested for two determinants of reference points. For the case of a habit formation reference point, the findings showed loss aversion but suggested model misspecification. In the case of a relative income reference point, the findings were inconsistent with reference-dependent preferences and loss aversion, instead indicating a standard concave utility function. In neither case was there a discontinuity in utility at the reference point.

Further analysis split GHQ into positive and negative affect to test whether negative affect was more affected by loss aversion, in line with the psychology literature where losses are considered qualitatively distinct from gains (Argyle, 1987; Rajabi and Sheykhshabani, 2009; Bayliss et al., 2014). This hypothesis was not supported; instead it seems that food consumption in general has a larger impact on negative than positive well-being. Future research could explore different theoretical specifications, perhaps testing more complex determinants of reference points such as rational expectations (Koszegi and Rabin, 2006, 2007, 2009) which was not permitted by the present data. Preference parameters could also be used to simulate policy and welfare analysis, as Finkelstein et al. (2012) do for state-dependence in health.

Overall, the results were mixed in their support of models of reference-dependent preferences and loss aversion. However, this kind of data will definitely be useful for investigating the structure of utility functions, particularly ‘behavioural’ varieties, in the future.

Chapter 3

Living in Your Parents' Shadow: Parents' Earnings as an Aspiration for Children

1 Introduction

Intergenerational mobility is often the focus of public policy and is likely to change from generation to generation. This paper provides the most up to date estimates of intergenerational mobility in England, the first for the so-called 'millennial' cohort. As well as estimating the aggregate Intergenerational Earnings Elasticity (IEE), it uses quantile regression of children's earnings on linear bins of parents' earnings to isolate the conditional IEEs at different points in the distribution. Results show that the aggregate IEE is low but likely downward biased due to life cycle effects, and that conditional IEEs increase with parents' earnings but decrease with children's earnings. As this pattern is not easily explained by conventional human capital models, a model is constructed where parents' earnings act as an aspiration, or pseudo-reference point, for children. In the model children invest in human capital in order to increase their earnings to this aspirational level which implies that children above (below) their parents' earnings will have lower (higher) IEEs, producing the pattern observed in the data.

This paper contributes to three strands of literature. Firstly, it contributes to the empirical literature on intergenerational mobility by providing up to date IEE estimates for England across the earnings distribution (Dearden et al., 1997; Grawe, 2004; Chetty et al., 2014). These are the first estimates of contemporary mobility in England because the final wave of the dataset used, the Longitudinal Study of Young People in England

(LSYPE), has only recently become available and is the only contemporary dataset on mobility in England. Furthermore, this paper is the first to use Grawe (2004)'s quantile regression method on English data to investigate how mobility varies at different points across the joint distribution of parents' and children's earnings.

Secondly, by proposing aspirations as an explanation for the observed pattern of IEEs this paper contributes to the literature on the causes of intergenerational mobility, which has looked at a number of mechanisms including credit constraints, early childhood investment and genetics but has not studied aspirations empirically (see Black and Devereux (2011) for an overview). Common approaches are to use detailed data – for instance on biological versus adopted parents or identical versus non-identical twins – to parse out environmental versus genetic influences. Alternatively, sudden policy changes which affect parents' earnings, such as plant closures or union status, can provide exogenous variation in parents' earnings (Shea, 2000). A final approach is to use proxies such as IQ for ability in order to control for omitted variables correlated with both children's and parents' earnings. The quantile regression approach in this paper is less widely used but can speak to multiple potential mechanisms at once, and has relatively low data demands.

Thirdly, by proposing a novel test for the effects of an unobserved reference point (in the form of an aspiration), the paper contributes to the empirical literature on reference points (Kahneman and Tversky, 1979; Camerer, 2000; Barberis, 2013) and in particular on reference points as aspirations (Blackburn and Chivers, 2015; Dalton et al., 2016; Besley, 2017; Genicot and Ray, 2017). Field tests of reference point models are becoming more common, but they typically focus on case studies with a salient reference point, a unique set of policy changes that affect this reference point, and an unusually detailed dataset measuring all of the relevant variables. Two such recent examples are retirement (Seibold, 2017) and unemployment benefits (Dellavigna et al., 2017), both of which are hypothesised to act as reference points affecting labour supply. This paper instead uses a standard cohort dataset and asks whether we can still observe the effects of reference points (as aspirations) on lifetime earnings in an intergenerational setting, making the method used potentially more general.

I use a contemporary English cohort dataset, the Longitudinal Study of Young People in England (LSYPE) (Centre for Longitudinal Studies, 2018), which follows children born in 1989-1990 and their parents. I estimate the aggregate log-log IEE, which measures the proportional change in children's earnings from a change in parents' earnings. Results show that a 1% increase in parents' earnings is associated with a

0.19% increase in children's earnings. I also estimate the rank-rank IEE, which measures the effect of parents' position in the earnings distribution on children's position in the earnings distribution. Results show that when both parents and children are ordered from 1 to n in their respective distributions, if the parents move up one rank it increases the child's rank by an average of 0.26.

Both estimates are lower than either estimates for previous cohorts in the UK (Dearden et al., 1997) or current cohorts in the US (Chetty et al., 2014), though the exact measure and estimation strategy is not the same across studies. As children are observed at age 27 in the LSYPE, it is also likely the present estimates are downward biased by life-cycle effects. Individuals with different levels of lifetime earnings tend to have different shaped earnings functions across their life cycle. Even after adjusting for age this can reduce estimated earnings persistence, as shown in a simple framework by Grawe (2006). Thus, the estimates are unlikely to be accurate measures of the true value of lifetime earnings persistence in England.

IEEs across the joint distribution of children's and parents' earnings give a fuller picture of mobility and make biases easier to account for, helping to discern which mechanisms are likely to be driving mobility – which is essential for policies aiming to improve it. Following Grawe (2004), I use quantile regression of children's earnings on linear bins of parents' earnings to show how intergenerational mobility varies with both children's and parents' earnings. The estimated coefficients are interpreted as the IEEs of a given quantile q of children's earnings with respect to a given bin j of parents' earnings. Grawe uses Canadian data and only interprets his results to reject the hypothesis of credit constraints without proposing an alternative explanation. I show that my results are also inconsistent with credit constraints, but I propose the aspirations model to explain the observed pattern of IEEs.

The quantile approach provides a closer look at intergenerational mobility than standard aggregate estimates of earnings persistence, and it differs from the transition matrix used by authors such as Dearden et al. (1997); Chetty et al. (2014) in two main respects. Firstly, as it does not use regression the transition matrix approach does not tell us the degree of variation in children's earnings explained by parents' earnings, nor does it allow us to control for confounding factors. Secondly, the transition matrix estimates the probability a child born into a family with (equivalised) earnings at a certain percentile will end up another percentile – say, the probability a child born into the 10th percentile will earn in the 90th percentile when they are older. The quantile approach instead estimates the impact of a rise in parental earnings from a bin j on the

earnings of their child who is *already in* earnings quantile q . This is best interpreted as a counterfactual e.g. if a family at the 90th percentile had earned even more, by how much would their child at the 10th percentile's earnings have increased — and would this differ if their child were at the 50th, or 90th percentiles? Conditional IEEs give a fuller picture of how family earnings are transmitted to children, for given levels of children's and parents' earnings.

Results show that conditional on parents' earnings, IEEs are lower for higher quantiles of children's earnings, falling by about 0.2 between the 10th – 90th quantiles. Conversely, conditional on children's earnings IEEs are higher for higher ranges of parents' earnings, rising by about 0.1 from the lowest to the highest bin. Further analysis shows this result is qualitatively insensitive to different measures of earnings, the number of splines, or other methodological choices. It is argued that this pattern cannot be explained by standard human capital models, drawing particularly from Han and Mulligan (2001), who explore the implications of heterogeneities in ability and intergenerational altruism for mobility.

Consequently I develop a new, two-period model where children aspire to earn as much as their parents, investing in human capital to reach this aspirational level. In line with recent literature on aspirations (Blackburn and Chivers, 2015; Dalton et al., 2016; Besley, 2017; Genicot and Ray, 2017) I model aspirations as a reference point. This is incorporated into a model where children choose between consumption now and investing in human capital to increase their earnings to reach their aspiration in the next period. The aspiration is an increasing function of parents' earnings but unobserved at the individual level. Children experience loss aversion – in the form of a penalty to their utility – when their earnings are below this aspiration, so they will invest in human capital at a higher rate to reduce their losses. The model predicts that the IEE will be higher for children who are below their aspiration and lower for children who are above their aspiration.

It is argued that this model accommodates the observed pattern of IEEs across the joint distribution of children's and parents' earnings. The intuition is that if a child comes from a low earning family background but has high earnings, it's likely their earnings are above their aspirations. This means they will have less incentive to accumulate human capital and so will have a lower IEE. In general, for a given level of parents' earnings high earning children will be more likely to be above their aspirations than low earning children and so will have a lower IEE. Similarly, for a given level of children's earnings those with low earning parents will be more likely to be above their

aspiration than children with higher earning parents, and so will also accumulate less human capital and have a lower IEE. This is precisely the pattern we observe in the data and formal tests verify that the results are consistent with the model.

The framework proposed in this paper has a number of appealing features as a method for detecting the effects of unobserved aspirations (or any reference point). Firstly, it requires mild assumptions at the individual level — in this case, only that aspirations are monotonically increasing in parental earnings. Secondly, it is flexible and could feasibly accomodate a variety of influences on aspirations, for instance region or peer effects. Thirdly, it is easy to interpret and straightforward for researchers to implement. Fourthly, the data demands are relatively low: earnings data from standard cohort datasets is all that is required. Finally, it uses a revealed preference approach, which serves as a (complementary) alternative to the more common stated preference approaches to aspirations and other ‘psychological’ variables such as beliefs (Jensen, 2010; Laajaj, 2017; Boneva and Rauh, 2018).

The chapter proceeds as follows. Section 2 outlines the dataset, the Longitudinal Study of Young People in England (LSYPE), including the methods used to make parents’ and children’s earnings comparable. Section 3 investigates intergenerational mobility in contemporary England, estimating IEEs across the distribution. Section 4 builds the aspirations model and shows how it explains the observed pattern, testing its implications formally. Section 5 concludes and discusses implications for future work.

2 Data and Earnings Derivation

This section summarises data from the Longitudinal Study of Young People in England (LSYPE), a nationally representative cohort study for contemporary England (Centre for Longitudinal Studies, 2018). The LSYPE covered the period from 2004-2016 and followed the families of children born between 1989-1991, gathering data on parents’ earnings, family background and demographics, as well as the childrens’ adult earnings in the final wave in 2015/16. The LSYPE has been used to measure intergenerational mobility in education in England (Anders, 2012; Blanden and Macmillan, 2016) but at the time of writing the data from the final wave has not yet been used to estimate mobility in earnings. Table 3.1 shows a set of demographic variables for the children in the LSYPE, while Table 3.2 shows similar variables for their parents.

Gross household parental earnings is measured in the first 4 waves of the LSYPE from 2004-2007. Earnings are reported as banded, so the mid point of each band is

Table 3.1: Demographic variables in the LSYPE (Children)

	Mean	Standard Deviation
Panel A: Standard Variables		
Number of children	0.21	0.58
Subjective health status (1-5 scale)	2.14	1.18
Self-Reported Preferences (1-10 Scale)		
How trusting are you?	6.20	2.49
How often do you take risks?	5.90	2.34
How patient are you?	6.04	2.63
Panel B: Category Variables		
Male	0.44	0.50
White	0.68	0.47
Marital Status		
Living as couple	0.11	0.31
Married	0.53	0.50
Widowed	0.08	0.27
Divorced	0.05	0.22
Qualifications		
Degree or higher	0.21	0.41
No qualifications	0.40	0.49
Employment Status		
Self-employed	0.07	0.25
Employed	0.75	0.44
Unemployed	0.06	0.23
Student	0.05	0.22
Long term sick/disabled	0.02	0.13
Region		
London & Southeast	0.29	0.45
North	0.38	0.48
Midlands & East	0.27	0.44

The means and standard deviations of variables measuring the economic and demographic characteristics of the children in the LSYPE from 2004-2016 in the final sample ($n = 5,177$). Panel A variables are positive values, while Panel B variables are measured in proportions between 0 and 1. All given moments exclude missing values for that variable.


Table 3.2: Demographic variables in the LSYPE (Parents)

	Mean	Standard Deviation
Panel A: Standard Variables		
Age (main)	41.15	8.88
Age (second)	43.94	6.39
Household size	4.23	1.40
Panel B: Category Variables		
Male (main)	0.28	0.45
White (main)	0.72	0.45
White (second)	0.72	0.45
Marital Status		
Single	0.08	0.27
Living together	0.07	0.25
Married	0.66	0.47
Divorced	0.10	0.30
Qualifications (Main)		
Degree or higher	0.12	0.32
No qualifications	0.25	0.43
Employment Status (Main)		
Self-employed	0.07	0.25
Employed	0.67	0.47
Unemployed	0.06	0.23
Student	0.01	0.11
Homemaker	0.24	0.42
Sick/Disabled	0.03	0.18

The means and standard deviations of variables measuring the parents of the children in the LSYPE from 2004-2008 in the LSYPE from 2004-2016 in the final sample ($n = 5,776$). Panel A variables are positive values, while Panel B variables are measured in proportions between 0 and 1. All given moments exclude missing values for that variable.

used with twice the minimum for the top band¹, a method which has been shown not to adversely affect estimation and inference (Fintel, 2006). Parents who are above 60 or below 20 are dropped to keep the sample to standard working-age, although as Figure 3.1 illustrates most parents are between the ages of 30 and 50 and as such the exclusion does not affect the results. The estimated average annual parental earnings in the sample (household equivalised) is £18,297, a similar estimate to Anders (2012) in the same dataset, albeit slightly smaller as his analysis uses weights whereas the present analysis does not. As the LSYPE aims to measure all of the resources avail-

Figure 3.1: Age Distribution



Distribution of the age of the child's self-reported main parent in the LSYPE ($n = 5,776$).

able to a child, only total household earnings are measured. Some of the recent literature has considered this a more appropriate measure for contemporary mobility than conventional father-son and mother-daughter pairings due to the entry of women into the labour force and the rise of single parent families (Belfield et al., 2017). Throughout this paper, both parents' and childrens' household earnings are equivalised by the

¹As the latter is a small part of the distribution, this choice does not affect any of the results in this paper

square root of the number of people in the household, which measures the child's effective standard of living growing up. This dominates other variables as a measure for a child's access to resources.

However, in the context of the model developed later – where parents' earnings act as an aspiration for children – it is less clear this measure dominates other measures. On the one hand, if two households have the same total earnings but in one household an only child enjoys the lion's share of resources; whereas in the other four siblings have to share the same amount, it is feasible the aspirations of the children will differ between households. On the other hand, it could be argued that a household with one parent earning £40,000 would result in children having different aspirations to a household with two parents each earning £20,000, which equalisation does not take into account. However, the current dataset does not permit us to test the latter possibility, so it is assumed that children's aspirations depend on their effective standard of living growing up.


Of most interest for intergenerational mobility is earnings persistence, or what in this paper is termed the Intergenerational Earnings Elasticity (IEE). Typically IEEs are calculated using the permanent earnings of both parents and children but the LSYPE contains only one observation for children's earnings — in the final wave, when they are approximately 27 years old. To make the parents' and children's earnings comparable we can predict parents' earnings at the age of 27 by running the following life cycle regression:

$$y_{it}^p = \alpha_i + \beta_1 age_{it} + \beta_2 age_{it}^2 + \beta_3 age_{it}^3 + h(t) + v_{it} \quad (3.1)$$

Where y_{it}^p denotes parents i 's earnings in year t , with a sample of n parents and T years in total, and α_i is an individual fixed effect. Unfortunately as shown in Figure 3.1, the age range at which parents are observed is quite narrow in the LSYPE, and moreover individuals are only observed for a four year period. So, while it is possible to estimate a parametric life-cycle earnings equation from such data, more precise estimates of the important shape parameters can be obtained from a larger sample with a broader age range and more observations per individual covering a longer time period.

In order to retrieve precise estimates of $\beta_1, \beta_2, \beta_3$, the estimation of (3.1) is instead run in the British Household Panel Survey/Understanding Society (BHPS/USoc) from the previous chapter, which ran from 1991-2016 and so covers a similar time period but is much larger, consisting of about 610,000 observations in total, and as shown by 3.2 the distribution of age is much broader in this dataset.

Figure 3.2: Age Distribution



Distribution of age in the combined BHPS/USoc Dataset ($n = 612, 771$).

Table 3.3 shows the results from regression (3.1) in the BHPS/USoc using only ages 20-60. All age coefficients are strongly significant and follow the expected life-cycle ‘hump-shaped’ pattern, increasing during youth and peaking in middle age, then tailing off around retirement. As Figure 3.3 shows, the parents’ estimated life cycle function varies by around 30% over the relevant age range, so this age adjustment is warranted. We can then take the estimated coefficients of the shape of the age profile and time effects, assuming they are the same for everyone, and subtract them from our raw earnings variable in the LSYPE to obtain an age and time adjusted measure \tilde{y}_{it}^p :

$$\tilde{y}_{it}^p = y_{it}^p - \hat{\beta}_1 age_{it} - \hat{\beta}_2 age_{it}^2 - \hat{\beta}_3 age_{it}^3 - \hat{h}(t) \quad (3.2)$$

The final step is to estimate the parents’ earnings at age 27, using the BHPS coefficients again:

$$y_{it}^p = \tilde{y}_{it}^p + \hat{\beta}_1 27 + \hat{\beta}_2 27^2 + \hat{\beta}_3 27^3 \quad (3.3)$$

Table 3.3: Life-Cycle Profile of Earnings (BHPS/USoc Data)


Parents' Earnings (log)	
(1)	
Age	0.0715*** (0.00244)
Age Squared	-0.00114*** (0.0000500)
Age Cubed	0.00000658*** (0.000000316)
Observations	603501

a. This table shows regressions of log parents' earnings on a cubic for age in the combined BHPS/USoc data.

b. Robust standard errors in parentheses.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.3: Parents' Life-Cycle Function

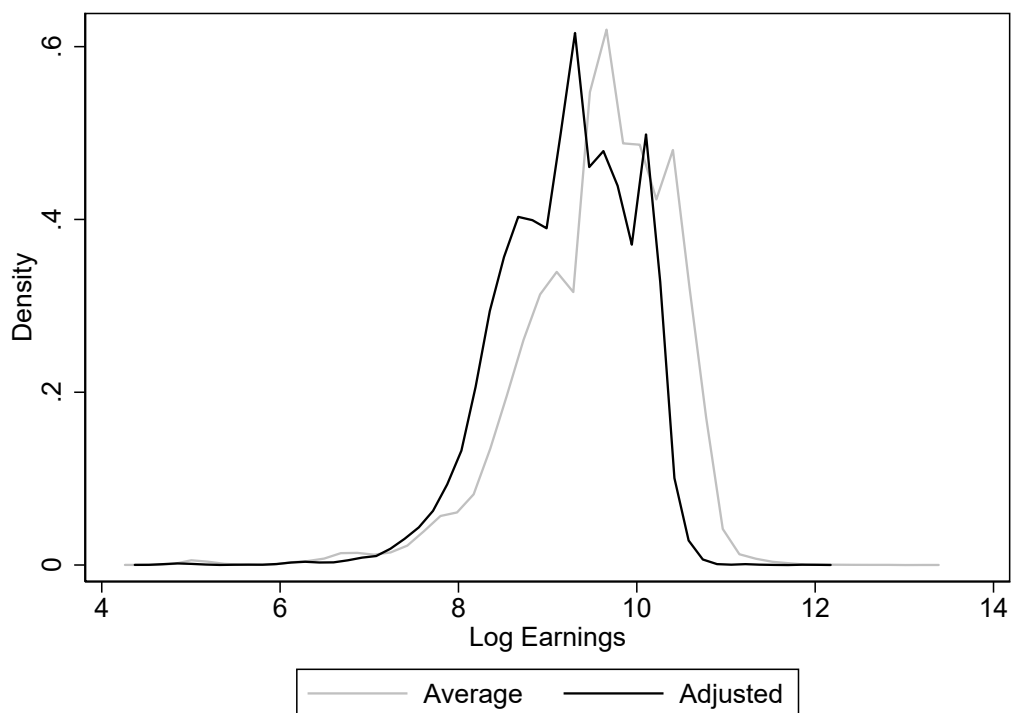


Logged earnings as a function of parents' age using the estimates in Table 3.3.

Standard errors for the life-cycle estimates are bootstrapped throughout the procedure and are clustered by individual. The bias-correction procedure for non-normality of error terms in Stata is used, though it does not much affect the coefficients. The process actually yields up to four estimates per parent - though they are all similar - so the average of these is taken. The final sample contains 5,177 pairs of children's earnings and comparable, age-adjusted parents' earnings.

Figure 3.4 shows the difference between parents' earnings adjusted for age and time effects and the raw mean of parents' earnings across the period 2004-2007, and it is evident that the former distribution has a lower mean. This is expected because the average age of parents in the sample is 42, substantially higher than the age of the children. Nevertheless, the raw average will still be used as a robustness check for the tests of the aspirations model in the appendix.

Figure 3.4: Parents' Earnings, Average and Age-Adjusted



The distribution of parents' earnings when they are averaged from 2004-2008 (light grey density) versus the distribution when they are adjusted for life cycle effects (dark grey density).

3 Intergenerational Mobility in Contemporary England

This section uses the LSYPE to investigate mobility in England over the sample period. It reports transition matrices, both log-log and rank-rank OLS regressions, and IEEs estimated across the distribution using the quantile regression method. Table 3.4 shows the transition matrix using the child's observed earnings at 27 and the parents' estimated earnings at 27, split into quintiles. Perfect mobility would result in each cell being 0.2, representing a 20% chance of ending up in each quintile no matter which quintile you were born into. Conversely, perfect immobility would be an identity matrix, with all diagonals equal to 1 and every other cell equal to 0. As would be expected the data are somewhere inbetween. Many cells in the middle are close to 0.2, suggesting a reasonable degree of middle class churn, while those on the edges are consistent with low mobility at the top and bottom. This result is in line with other findings in the U.K. (Dearden et al., 1997) and also the U.S. (Chetty et al., 2014).

In the sample, a child born into the top quintile has over a 30% chance of ending up in the same quintile and only a 14% chance of ending up in the bottom quintile. Conversely, a child born into the bottom quintile has only a 10% chance of ending up in the top quintile, but a 29% chance of ending up back in the bottom quintile. On the other hand, a child born into the middle quintile has just over a 20% chance of ending up within one quintile of their parents, and just under a 20% of ending up at either of the extreme quintiles — close to perfect mobility in the middle.

Table 3.4: Transition Matrix

Elasticity	Parent's Quintile				
Child's Quintile	1 st	2 nd	3 rd	4 th	5 th
1 st	0.293	0.247	0.193	0.168	0.099
2 nd	0.223	0.216	0.207	0.183	0.172
3 rd	0.187	0.214	0.215	0.210	0.174
4 th	0.157	0.176	0.201	0.226	0.240
5 th	0.141	0.146	0.185	0.214	0.315

Probability of a child with parents in a particular quintile having earnings in each quintile when they are older. The columns are the parents' earnings quintile while the rows are the child's. Earnings are taken from the LSYPE and children's are measured in 2015/16 at age 27. Parents' earnings are measured from 2004-2008, adjusted for life cycle effects to yield a predicted value for the parents aged 27.

Table 3.5 shows the standard OLS estimate of IEEs – aggregate earnings persistence – both in log-log and rank-rank earnings regressions. The univariate log-log estimate in column (1) implies a 1% rise in parents' earnings is associated with a 0.19% rise in

children's earnings. As column (3) shows the rank-rank estimate is higher, implying that an increase in the parents' rank of one increases the child's rank by 0.26 points. This is consistent with the point made in Chetty et al. (2014) that rank measures allow zeroes to be coded and that this increases estimates of earnings persistence.

The log-log estimate is at the smaller end of previous estimates for the UK in the literature, which range from a lower bound of 0.2 for crude, unadjusted OLS (which is likely to be downward biased, see below) up to an upper bound of 0.7 IV estimates — which are likely upward biased because most instruments in the literature are invalid (Dearden et al., 1997). At 0.16, the univariate rank-rank estimate is lower than Belfield et al. (2017)'s estimate of 0.22 for the Baby Boomer cohort and that of 0.32 for Generation X. It is also low in comparison to Chetty et al. (2014)'s contemporary estimate of about 0.34 for the U.S. However, it should be noted that both the household equivalised measure and the age-adjustment procedure in this paper are unique². Therefore estimates are not directly comparable to previous literature.

As Grawe (2006) demonstrates using a simple framework of lifetime earnings with both permanent as transitory components, there is good reason to believe that a steeper lifecycle earnings profile of children relative to parents will result in a downward bias in the IEE (this applies to both log and rank estimates), even after it is adjusted for age effects to eliminate transitory shocks. The divergence between the permanent and observed earnings of children relative to that for parents reduces the estimated coefficient. In the LSYPE parents are old relative to children, and there was an expansion of higher education in the UK – which is associated with a steeper lifecycle profile – over the sample period. These factors make it likely the earnings of children we are observing are a much smaller fraction of their annualised lifetime earnings than those of their parents, explaining the low estimates.

In column (2) the addition of covariates creates a tangible but limited decline in the estimated log-log coefficient. One might expect a bigger decline as covariates include variables which would be expected to capture some of the differences in human capital between children such as health, education, and occupation. Even more surprising is that there is no statistically significant difference in the rank-rank coefficient after the addition of covariates. These results imply that parental earnings explain children's earnings beyond these variables, though there are the notable omissions of childhood investment, grades in education, and intangible human capital (such as job-specific

²Belfield et al. (2017) do adjust for household size in a similar manner but as they observe children at much later ages, they do not perform the same age-adjustment as the present study.

Table 3.5: Earnings Persistence

	Log Children's Earnings		Rank Children's Earnings	
	(1)	(2)	(3)	(4)
Log Parents' Earnings	0.192*** (0.0173)	0.153*** (0.0174)		
Rank Parents' Earnings			0.163*** (0.0104)	0.123*** (0.0106)
Observations	5177	5177	5177	5177

a. This table shows regressions of children's earnings on parents' earnings in the LSYPE. Columns (1) and (2) use log earnings while columns (3) and (4) use the rank of earnings, where both parents and children are ordered from 1 to n in their respective distributions. As the rank includes zeros, the sample size is slightly higher. Covariates include race, gender, measured preferences, health, education, and occupation.

b. Robust standard errors in parentheses.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

skills).

Table 3.8 shows the results of the regression using log-log estimates only and including the coefficients for covariates. Rank-rank is not used because it makes the coefficients on the covariates harder to interpret. Covariates are split into two different categories: column (2) shows only 'exogenous' covariates which are beyond the child's control such as race, gender and preferences (assumed fixed in the neoclassical sense). Column (3) adds covariates which could be considered 'endogenous' variables as they measure accrued human capital: health, employment status, and education. The rationale is that the two have different interpretations from the perspective of individuals decisions and for policy. Higher earnings due to being white and/or male are out of the individual's control whereas pursuing education, leading a healthy style and seeking employment are not.

Being male increases earnings by 14% in the fullest specification in column (3) and whites earn almost 7% more. Higher levels of self-reported willingness to trust increase earnings, while risk aversion has no impact. Curiously, more patient people tend to earn less, which is at odds with standard human capital models but could be a result of observing children so early in the life cycle, similarly to the argument in Grawe (2006) referenced above³. Healthy people earn 17% more than unhealthy people and those with degrees earn approximately 27% more than those without. In both these cases the effect may be due to unobserved heterogeneity because more able,

³If patient people are willing to wait for higher earnings when they are older, they will have steeper life cycles and this could explain the effect.

Table 3.6: Earnings Persistence with Covariates

	No Covariates (1)	Exogenous Covariates (2)	Endogenous Covariates (3)
Log Parents' Earnings	0.192*** (0.0173)	0.182*** (0.0179)	0.153*** (0.0174)
Trust		0.0220*** (0.00568)	0.0142*** (0.00542)
Risk		0.00688 (0.00626)	0.0102* (0.00599)
Patience		-0.0189*** (0.00487)	-0.0166*** (0.00467)
Male		0.150*** (0.0234)	0.141*** (0.0227)
White		0.0163 (0.0276)	0.0680** (0.0282)
Full Time Employment			1.048*** (0.0898)
Healthy			0.166*** (0.0403)
Degree			0.274*** (0.0252)
Observations	5177	5177	5177

a. This table shows regressions of children's earnings on parents' earnings in the LSYPE. Columns (1) and (2) use log earnings while columns (3) and (4) use the rank of earnings, where both parents and children are ordered from 1 to n in their respective distributions. As the rank includes zeros, the sample size is slightly higher. Covariates include race, gender, measured preferences, health, education, and occupation.

b. Robust standard errors in parentheses.

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

motivated and fastidious people will be healthier and more qualified. As would be expected employment has a big impact on earnings, doubling them relative to other employment categories.

3.1 Elasticities Across the Distribution

In order to better understand the determinants of intergenerational mobility, this paper follows Grawe (2004) by estimating the pattern of IEEs across the distribution. Different theories may result in the same aggregate earnings persistence but different patterns across the distribution, as Grawe points out in the case of credit constraints. Such estimates can therefore be used to distinguish between different mechanisms in a way that aggregate IEEs cannot. It could be argued that estimating a wide range of IEEs simply compounds the problem of bias, since we now have many coefficients to worry about instead of one. However, biases which are ambiguous in the aggregate may be unidirectional at particular points in the distribution. This is the case for many of the theoretical examples in Han and Mulligan (2001) and also for the aspirations model in this chapter, both of which are discussed in more detail later.

This section estimates IEEs across the joint distribution of parents' and children's earnings using quantile regression of y^c on linear splines of y^p (both in logs throughout), shown in regression (3.4). These estimates should not be termed 'earnings persistence' which is why this chapter uses the term 'Intergenerational Earnings Elasticity' (IEE), which is easier to interpret across the joint distribution. These estimates capture how each quantile q of childrens' earnings depend on parents' earnings, conditional on the parent being in a particular earnings range, as captured by bin j . This gives a richer picture of earnings mobility in contemporary England than the aggregate regressions reported in Tables 3.5 & 3.8 while also moving beyond the transition matrix approach to analysing the joint distribution in Table 3.4.

The specification is:

$$y_q^c = \theta_{0,q} + \theta_{1,q} \cdot \mathbb{1}(y_i^p < s_1) \cdot y_i^p + \sum_{j=2}^J (\theta_{j,q} \cdot \mathbb{1}(s_{j-1} \leq y_i^p < s_j) \cdot y_i^p) + \theta_{J,q} \cdot \mathbb{1}(y_i^p \geq s_J) \cdot y_i^p + \delta_q X_i + \varepsilon_{i,q} \quad (3.4)$$

Covariates are collected in the vector X_i and are included in some specifications. Each linear bin of y_i^p is cut off at the 'knots' s_j , with $j = 1, \dots, J$ and a total of J bins. These bins split y^p into continuous ranges with cutoff points so that the coefficient $\theta_{j,q}$, the

linear spline function, can vary depending on the level of y^p . Bins are coded so the splines (or IEEs) do not join at the knots s_j . Although we might expect children's earnings to vary continuously with parents' earnings, this specification means the IEEs can be interpreted as the elasticity for a subgroup of children earning at quantile q with parents earning in bin j . It is an average of the different slopes of individuals rather than an estimate of a single underlying function, and so may not join with other subgroups if the earnings function varies across the distribution. This is precisely the case in the context of the model developed later, where children may have different earnings slopes depending on their aspirational level of earnings. A specification where splines join at the knots, so that the function is continuous, is explored in the Appendix as a robustness check.

Table 3.7 shows the results with $J = 3$ and the cutoffs s_j simply set at the terciles of parents' earnings, so that each spline contains the same number of data points. The IEE for the lowest bin of parental earnings $\hat{\theta}_{1,q}$ at the 25th quantile is 0.147, implying a 1% increase in the earnings of low-earning parents is associated with a 0.15% increase in the earnings of low-earning children. The IEE decreases to 0.109 for the 50th quantile and 0.096 for the 75th quantile. For the next bin of parents, captured by $\hat{\theta}_{2,q}$, the IEEs are higher across all 3 quantiles: a 1% increase in the earnings of medium-earning parents increases children's earnings at the 25th quantile by 0.15%, by 0.11% at the 50th quantile, and by 0.9% at the 75th quantile.

In general, for a given quantile q the IEEs increase as the bin j of parents' earnings increases so that higher earning parents have a bigger effect on their children's earnings across the distribution. Furthermore, for a given bin j the IEE falls as q rises — for example $\hat{\theta}_{3,q}$ falls by 0.057 between the 25th and 75th quantiles. This shows that parents' earnings have less of an effect on the earnings of higher earning children.

Covariates are now added to the quantile regression to test (a) whether the covariates themselves have varying effects across the distribution and (b) whether they change the effect of parents' earnings observed in Table 3.7 across the distribution⁴. Table ?? shows the estimated IEEs with all covariates included in the regression and the IEEs are reduced across the board as would be expected: similarly to the standard OLS results in Table 3.8, most of the coefficients fall by 0.04-0.06. In the case of $\hat{\theta}_{2,50}$ the coefficient becomes statistically insignificant where it was previously significant.

⁴It should be borne in mind that the more covariates are added the more difficulty *Stata* has making the quantile estimates to converge, which is likely a consequence of a fairly low sample size. With more covariates, convergence fails more often so more attempts are necessary to retrieve a sufficient number of estimates.

Table 3.7: IEEs across the distribution (unjoined splines, without covariates)

		25th Quantile	50th Quantile	75th Quantile
		(1)	(2)	(3)
IEEs	$\hat{\theta}_{1,q}$	0.147*** (0.0428)	0.109*** (0.0387)	0.096*** (0.0354)
	$\hat{\theta}_{2,q}$	0.153*** (0.0380)	0.113*** (0.0353)	0.094*** (0.0313)
	$\hat{\theta}_{3,q}$	0.162*** (0.0354)	0.126*** (0.0334)	0.105*** (0.0331)
Observations		5177	5177	5177

a. This table shows the elasticities from quantile regressions of children's earnings on bins of parents' earnings in the LSYPE as in regression (3.4). Each column shows a different quantile of children's earnings while each row shows a different bin of parents' earnings, with $\hat{\theta}_{1,q}$ representing the lowest and $\hat{\theta}_{3,q}$ representing the highest bin values.

b. Bootstrapped standard errors in parentheses

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Despite this, the pattern observed is broadly the same as in Table 3.7: as q rises the IEEs fall; as j rises the IEEs increase.

Table 3.7 also shows the effect of the covariates themselves at each quantile. Men have a greater advantage at the top of the distribution, while for whites the effect rises in the middle and then falls at the top. The effect of trust is positive at both the bottom and top but insignificant in the middle, while risk aversion has a positive effect at the bottom only. Patience has a negative effect across the distribution which is highest in magnitude at the bottom and lowest in the middle. Employment has a larger effect at the bottom and at the top, implying it is more important for the poor, which is an intuitive result. Good health and having a degree also have smaller (though still positive and significant) effects at the 75th quantile, showing that human capital accumulation compresses the distribution. This is in line with Eide and Showalter (1998), who find the same for education on U.S. data.

A unifying interpretation of the results in Tables 3.7 is that the IEE depends on the *relative* distance between parents' and children's earnings. When parents' earnings are higher than children's, the IEE is high; when parents' earnings are lower than children's, the IEE is low. Thus for a given bin of parents' earnings j , as the quantile of children's earnings q increases, the IEE falls. Conversely, for a given quantile q , as the range captured by bin j of parents' earnings increases, the IEE rises. This

Table 3.8: Earnings Persistence with Covariates


	25th Quantile (1)	50th Quantile (2)	75th Quantile (3)
$\hat{\theta}_1$	0.117*** (0.0396)	0.0638* (0.0372)	0.0787*** (0.0298)
$\hat{\theta}_2$	0.121*** (0.0362)	0.0693** (0.0349)	0.0785*** (0.0267)
$\hat{\theta}_3$	0.129*** (0.0331)	0.0794** (0.0320)	0.0853*** (0.0254)
Trust	0.0129** (0.00518)	0.0118 (0.00739)	0.00800** (0.00326)
Risk	0.0201*** (0.00712)	0.00747 (0.00508)	0.00240 (0.00463)
Patience	-0.0205*** (0.00423)	-0.0133*** (0.00377)	-0.0168*** (0.00290)
Male	0.134*** (0.0236)	0.135*** (0.0234)	0.174*** (0.0174)
White	0.100*** (0.0209)	0.165*** (0.0244)	0.0958*** (0.0260)
Full Time Employment	1.210*** (0.0511)	0.939*** (0.0858)	0.764*** (0.106)
Healthy	0.167*** (0.0412)	0.166*** (0.0523)	0.133*** (0.0412)
Degree	0.287*** (0.0270)	0.236*** (0.0192)	0.237*** (0.0221)
Observations	5177	5177	5177

a. This table shows the elasticities from quantile regressions of children's earnings on bins of parents' earnings in the LSYPE as in regression (3.4). Each column shows a different quantile of children's earnings while each row shows a different bin of parents' earnings, with $\hat{\theta}_{1,q}$ representing the lowest and $\hat{\theta}_{3,q}$ representing the highest bin values.

b. Bootstrapped standard errors in parentheses

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.5: Pattern of IEEs



The pattern of IEEs across the joint earnings distribution of parents' earnings, on the x-axis, and children's earnings, on the y-axis (both in logs). The graph shows the IEEs – given by dashed grey lines – at four points in the distribution. At point (y_h^c, y_l^p) parents' earnings are low and children's are high, and the IEE is the lowest. At point (y_l^c, y_h^p) children's earnings are low and parents' earnings are high, and the IEE is the highest. At points (y_l^c, y_l^p) and (y_h^c, y_h^p) parents' and children's earnings are both low and high, respectively, and the IEEs take intermediate values.

pattern is shown in stylised form in Figure 3.5 for different pairs of (y_k^c, y_k^p) , where $k = l$ indicates a low value and $k = h$. The shallowest slope or lowest elasticity is at the point (y_h^c, y_l^p) , while the steepest is at (y_l^c, y_h^p) , with the other two elasticities taking intermediate values. The key intuition in the following model is that this is evidence that children whose earnings are low relative to their parents are experiencing 'loss aversion' and attempting to increase their earnings above those of their parents, resulting in the steeper slope of the IEE.

It is first important to establish that this pattern is not easily explained by the canonical Becker and Tomes (1979, 1986), where parents invest optimally in their children's human capital in accordance with innate parameters such as preferences or ability. While these parameters may vary across the distribution, none of the variants of the model imply the pattern we observe in the present study. Han and Mulligan (2001) explore a version of Becker-Tomes with heterogeneity in both altruism towards children and in innate ability. Heterogeneity in ability – assumed to be multiplicative with investment in human capital – induces an upward bias among high earning parents with high earning children, implying this group would have the highest IEE, which is not

what we observe in these results. Heterogeneity in altruism rates has a different implication: low earning parents with high earning children must have invested a higher proportion of their earnings into their children, implying they are more altruistic and have a higher IEE; while high earning parents with low earning children must be more selfish and have a lower IEE. Again this is contrary to what we see in the results above.

As noted by Becker and Tomes (1986), risk has ambiguous implications for their framework, depending on the chosen functional forms and the degree of uncertainty in the model. As a result, risk has received less systematic attention in the literature. For instance, in Han and Mulligan (2001) the parameter which governs risk aversion also governs the intertemporal substitution between children and parents, and its impact varies depending on the various iterations of their model. Thus, while the results above could be consistent with the presence of risk and heterogeneity in risk aversion, it is unclear there is an *ex ante* desirable set of assumptions about risk which would generate these results. To put it another way, while risk aversion might explain the above results, there is reason to believe it might explain any set of results with the right assumptions about functional form.

As Grawe (2004) demonstrates, introducing credit constraints into the Becker-Tomes framework means that persistence will be higher among low earning parents with high earning children, as these are the least able to borrow to fund their child's education. This implies higher IEEs at higher quantiles, which is neither what Grawe finds nor what we see in the results above⁵. A similar result follows from the Overlapping Generations model in Blackburn and Chivers (2015), where children invest in human capital with uncertain returns. Children have aspirations and experience loss aversion if their earnings are below this aspiration, leading them to refrain from human capital investment if there is a low chance they will reach their aspirations. As the authors themselves note, their model is likely to be observationally equivalent to credit constraints because it implies a high return among low earning children.

It is important that in the Becker-Tomes framework the parents make decisions, whereas the key intuition in this paper is that the children's decisions are generating this pattern of IEEs across the distribution, as shown in Figure 3.5. In particular, if parents' earnings act as an aspiration, children who are above (below) their parents' earnings have less (more) of an incentive to increase their earnings, thus producing a shallower (steeper) slope with respect to their parents' earnings. The next section develops a human capital framework where parental investment has already taken place

⁵Andrade et al. (2003) do find some evidence for this in Brazil, though their data are poor quality.

and children aspire to a level of earnings which depends on those of their parents. The children can invest in human capital themselves in order to increase their earnings from period to period and will do more so if their aspirations are higher, producing the pattern observed in the data. It should be borne in mind that this model is still in its early stages and has a number of limitations which will be addressed in the future.

4 Unobserved Aspirations Model

This section describes a simple human capital model of children who have aspirations which positively depend on their parents' earnings. The model is similar to the classic Becker and Tomes (1979, 1986) framework but here the child's decision to invest in human capital is made *after* they receive their initial draw of earnings rather than before. Parents' earnings have been determined in the past and so are taken as given by the child. We focus on the child's decision early in their working life, abstracting from the decisions of the parents. Although the parents' decision is important and could interact with the child's decision in interesting ways, this is the best way to generate straightforward predictions to interpret the empirical results above. Future work could look more into the theoretical mechanisms which arise when both decisions are taken into account.

There are n families and each child from family i has starting earnings in period 1, determined by the simple human capital equation $y_{1,i}^c = by_i^p + u_i$ with coefficient b and shock $u_i \sim N(0, \sigma^2)$. The interpretation is that through early childhood, school and other family influences – plus random luck – the parents invest in the child to give them a starting level of human capital, after which the child makes decisions about whether to invest more into human capital to increase their earnings further. Note that this can be thought of as the outcome of a Becker and Tomes (1979, 1986) framework even though the parents' decision is abstracted from for simplicity. From now on the i 's will be dropped for exposition purposes. The child faces a two-period problem similar to Genicot and Ray (2017), where they can invest in human capital in order to reach their aspirations in the next period, though here the investment is done by the child rather than by the parents. The child chooses optimally between the two periods: in period 1 they can consume their earnings c_1 or invest them into human capital I (eg education and training) in order to increase their earnings for the second period. Second period earnings are generated by the linear human capital function $y_2^c = I(1 + r) + y_1^c$, with a rate of return on human capital r . Second period utility is discounted by $\rho < 1$ and the

child solves:

$$\begin{aligned}
\max_{c_1, I} \quad & U(c_1, y_2^c) = \ln(c_1) + \rho(\ln(y_2^c) + \mu(\ln(y_2^c) - \ln(a))) \\
s.t. \quad & y_2^c = (1+r)I + y_1^c \\
& c_1 + I = y_1^c
\end{aligned} \tag{3.5}$$

The child receives increased utility when their second period earnings are at or above their aspirational level a , and decreased utility when their earnings are below it. Aspirations are assumed to be an increasing function of the parents' earnings $a = a(y^p)$, $a'(y^p) > 0$ and they enter the child's utility through $\mu(\cdot)$, a linear reference-dependent function similar to the previous chapters:

$$\mu(\ln(y_2^c) - \ln(a)) = \begin{cases} \gamma(\ln(y_2^c) - \ln(a)), & \text{if } y_2^c \geq a \\ \lambda(\ln(y_2^c) - \ln(a)), & \text{if } y_2^c < a \end{cases}$$

it is assumed that $\lambda > \gamma$ i.e. 'losses loom larger than gains': the child experiences loss aversion from not achieving their aspirations. The parents' investment determines the child's starting earnings while the parents' realised earnings determine the child's aspiration. The child does not know the function that determines the parents' earnings but they know their starting earnings y_1^c and observe their parents' earnings *ex post*. Crucially, the shock u may drive a wedge between the child's initial earnings level and their aspirational earnings, and the child will have to invest in human capital to make up the difference.

Maximising the child's utility function in (3.5) with respect to human capital investment I yields the following first order condition:

$$\begin{aligned}
\frac{dU}{dI} &= -\frac{1}{y_1^c - I} + \frac{\rho(1+r)}{I(1+r) + y_1^c} + \frac{\rho\mu'(1+r)}{I(1+r) + y_1^c} = 0 \\
&\iff \frac{1}{y_1^c - I} = \frac{\rho(1+\mu')(1+r)}{I(1+r) + y_1^c}
\end{aligned}$$

where μ' is the derivative of μ , so $\mu' = \gamma$ if $y_2^c \geq a$ and $\mu' = \lambda$ if $y_2^c < a$. Solving for I yields the optimal level of children's human capital investment I^* :

$$I^* = \frac{1 + \rho(1+\mu')(1+r)}{(1+r)(1 + (1+\mu')\rho)} y_1^c$$

substituting back in the equation for y_2^c gives us:

$$y_2^{c,*} = \frac{1 + \rho(1 + \mu')(1 + r)}{1 + (1 + \mu')\rho} y_1^c \quad (3.6)$$

We do not observe y_1^c in the data (we only have one wave of children's earnings) so (3.6) cannot be estimated, but since $y_1^c = by^p + u$ we can insert this back into (3.6) to re-express $y_2^{c,*}$ as a function of y^p :

$$y_2^{c,*} = \theta(by^p + u) \quad (3.7)$$

where $\theta = \frac{1 + \rho(1 + \mu')(1 + r)}{1 + (1 + \mu')\rho}$. The shock is orthogonal to y^p by assumption but as the term θ depends on μ' it is higher when y_2^c is below $a(y^p)$ and vice-versa. This induces a correlation between the shock and both y^p and y_2^c which could bias estimation, but as we will see later it does not affect the interpretation of the quantile regressions. Simple derivation shows that $\frac{d\theta}{d\mu'} = \frac{b\rho r}{(\rho(\mu'+1)+1)^2}$ is positive so that $\theta(\cdot)$ is increasing in μ' , which implies that it is higher for individuals below their aspirations, for whom μ' is higher.

Now consider the implications of the model for a standard cohort dataset such as the LSYPE, where we have a population of data points measuring child's and parents' earnings at a given point in their respective life cycles (y_i^c, y_i^p) , $i = 1, \dots, n$ with n families. As there is only a single cross-section, the time subscript for children's earnings is omitted. By simply inserting $\mu' = \gamma$ when $y_i^{c,*} \geq a(y^p)$ and $\mu' = \lambda$ when $y_i^{c,*} < a(y^p)$ into (3.7), we can make the following proposition:

Proposition 1: The true IEE for a given individual is:

$$y_i^c = \begin{cases} \theta_{\gamma,i} y_i^p + shock & y_i^c \geq a_i(y_i^p) \\ \theta_{\lambda,i} y_i^p + shock & y_i^c < a_i(y_i^p) \end{cases} \quad (3.8)$$

with $\theta_{\gamma,i} < \theta_{\lambda,i}$, i.e. the IEE will be higher for children who are below their aspirations. The result in (3.8) cannot be estimated directly for two reasons. Firstly, aspirations are unobservable at the individual level. Secondly, the function is conditional on the dependent variable. We must therefore take an indirect approach to teasing out the implications of the aspirations model for the data.

In line with the empirical results from the previous section, we are interested in how the aspirations framework predicts the variation in the effect of y_i^p on y_i^c – the IEE – across the distribution of y_i^p and y_i^c . According to the model a high y_i^c relative

to y_i^p means children at this point in the distribution are more likely to be *above* their aspirations, so on average their estimated IEE will be lower. Conversely, a high y_i^p relative to y_i^c means children are more likely to be *below* their aspirations, resulting in a lower estimated IEE. In other words: conditional on a value of y_i^c , the IEE increases with y_i^p . Conditional on a value of y_i^p , the IEE decreases with y_i^c , the same as the results above.

The logic of this prediction is in line with Figure 3.5 earlier. At point (y_h^c, y_l^p) children are likely to be above their aspirations so the slope of the IEE is shallow compared to the point (y_h^c, y_h^p) , where y_i^c is the same but y_i^p is higher. Similarly, at point (y_l^c, y_h^p) children are likely to be below their aspirations so the estimated IEE is higher than both point (y_h^c, y_h^p) , where y_i^c is higher, and point (y_l^c, y_l^p) , where y_i^p is lower. To further illustrate this pattern we can return to the quantile regression specification in regression 3.4, which allows us to make conditional estimates of the IEEs:

$$y_q^c = \theta_{0,q} + \theta_{1,q} \mathbb{1}(y^p \leq \bar{y}^p) * y^p + \theta_{2,q} \mathbb{1}(y^p > \bar{y}^p) * y^p + \epsilon_q \quad (3.9)$$

Since y_q^c is fixed using quantile regression, a will be higher relative to y_q^c for the upper splines of y^p and so a greater proportion of children will be below their aspirations. Thus, the estimated coefficients should be characterised by $\hat{\theta}_{2,q} > \hat{\theta}_{1,q} \quad \forall q$. The null hypothesis will be $H_0 : \hat{\theta}_{2,q} = \hat{\theta}_{1,q}$ against the one-sided alternative $H_A : \hat{\theta}_{2,q} > \hat{\theta}_{1,q}$. Furthermore, a higher quantile of y_q^c means the child is more likely to be *above* their aspirations and so as q rises $\hat{\theta}_{j,q}$ will fall for $j = 1, 2$. The null hypothesis to be tested is that for a given j , $\hat{\theta}_{j,\bar{q}} > \hat{\theta}_{j,q}$ for some $\bar{q} > q$. In the estimation, these will be set at $\bar{q} = 75$ and $q = 25$ so the test is interquartile. Regression (3.4) uses 2 bins of parental earnings, split at the median \bar{y}^p for simplicity, but the logic should hold for any number of splines provided there is sufficient data, and similar tests are run for 3 splines in the Appendix⁶.

A confounding force in this model is that children who have a high value of a will have more incentive to increase their final earnings, moving them up the distribution of y^c . This means that higher values of q could be associated with a *higher* IEE, contrary to the predictions above. Consider a simple example where a ‘child’ has first period earnings of $y_1^c = 10,000$ and can decide how much to invest in human capital

⁶Note also that the regression is in logs while the theory is in levels. This is because without taking logs regression (3.4) does not converge. But by taking logs the prediction in (3.8) will be eliminated since the parameters are absorbed into the constant. Thus the main results should not be thought of as a direct estimation of the aspirations model but a test of the basic intuition that it provides. In the future I would like to construct a more general model in which the prediction holds for logs.

to affect their earnings next period. Assume further that their parents' earnings were $y^p = 20,000$ and $b = 1$, so $u = 10,000$. Aspirations are $a = 20,000$ so the child is 10,000 short of their aspirations. Suppose further that the parameters ρ, r, λ are such that $\theta_\lambda = 2$. Following equation (3.6) $y_2^c = 20,000$, so the child's high aspirations have pushed them up the distribution of earnings by a substantial number and they would be captured by a higher quantile in regression (10).

However, we assume the return to human capital r is low year-to-year so this effect does not confound estimation. That is, we assume the initial earnings dominate the change in earnings that results from the child's human capital investment, so that the distribution of y^c is mostly retained between periods 1 and period 2. In fact, this is the most credible assumption: theoretically any value of r would have to be implausibly high to produce the magnitude of θ_λ shown in the previous example. In the previous example but with $r = 0.05, \rho = 1, \gamma = 1$ and $\lambda = 2$, $\theta_\lambda = 1.0375$ therefore $y_2^c = 10,375$. The intuition is that annual pay rises are usually a low proportion of starting salary, though over a lifetime this effect would certainly confound estimation.


Another issue was mentioned above: in the prediction of the model in (3.7) there is a correlation between y^c and the unobserved term θu . As θ is positive, the sign of this term depends on the shock to children's earnings u , which can be either positive or negative. The effect on standard OLS is therefore ambiguous, but in quantile regression the effect will vary across the distribution. To illustrate this, consider again the point (y_h^c, y_l^p) in Figure 3.5. Since y^c is high and y^p is low, we know that by assumption u is likely to be strongly positive and therefore the coefficient is biased upwards. Conversely, at (y_l^c, y_h^p) , u is likely to be negative and the coefficient biased downwards. But the prediction made was that the coefficients would be low at the former point and high at the latter point. Thus, the bias works in the opposite direction to the predictions and so if the pattern is still observed in the data it cannot be a result of the bias. Indeed, the actual estimates are a lower bound on the true effect of aspirations across the distribution.

4.1 Testing the Model

In order to test the predicted pattern more explicitly on the data, regression (3.4) is run – with the splines labelled h and l for 'high' and 'low' – and from the 10th to the 90th quantiles (in jumps of 10 to smooth out noise). Figures 3.6 and 3.7 show the point estimates of $\hat{\theta}_{h,q}$ and $\hat{\theta}_{l,q}$ (respectively) as the quantile of y^c increases. In both cases there is a steady decline as q increases as predicted, though for $\hat{\theta}_{l,q}$ the fall

is slower and the confidence intervals are larger, with the coefficient not statistically distinct from zero at some quantiles.

Figure 3.6: Quantile Coefficients ($\hat{\theta}_h$)




IEEs for high earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines.

Table 3.9 shows the results from an interquantile regression on Stata, which estimates the difference in given coefficients between two quantiles and shows whether or not it is statistically significant. Reported coefficients are the difference between the IEEs at the two quantiles. The chosen test is an interquantile test, which is robust to extreme values at the ends of the distribution - though the results are similar for the difference between the 10th and 90th, or 20th and 80th quantiles. Both splines decline by 0.079 between the two quartiles and the difference is significant at the 5% level.

Figure 3.8 shows both coefficients together but without confidence intervals for exposition purposes. By inspection it is clear that $\hat{\theta}_{l,q} < \hat{\theta}_{h,q}$ as predicted. Formally, the one-sided null hypothesis that the two coefficients are equal is rejected with 1% confidence at every quantile. Robustness checks – using an alternative earnings measure, different cutoffs for y^p , with covariates, with joined splines, and with 3 splines instead

Figure 3.7: Quantile Coefficients ($\hat{\theta}_l$)



IEEs for low earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines.

of 2 – are presented in the Appendix. All results are similar to the results presented here.

Table 3.9: Interquartile Difference in IEEs


Children's Earnings	
$\hat{\theta}_{l,75} - \hat{\theta}_{l,25}$	-0.079** (0.0371)
$\hat{\theta}_{h,75} - \hat{\theta}_{h,25}$	-0.079** (0.0325)
Observations	5177

a. This table shows the elasticities from interquartile regressions of children's earnings on bins of parents' earnings in the LSYPE. There are two bins of parents' earnings, with $\hat{\theta}_{l,q}$ representing the lowest and $\hat{\theta}_{h,q}$ representing the highest values. The estimated coefficients represent the difference between the effect of a given bin of parents' earnings on children's earnings at the 75th versus 25th quantiles.

b. Bootstrapped standard errors in parentheses

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.8: Quantile Coefficients ($\hat{\theta}_l$ & $\hat{\theta}_h$)



IEEs for both high and low earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The former is given by the solid grey line while the latter is given by the solid black line.

5 Conclusion

This paper has provided the most up-to-date estimates of intergenerational mobility in contemporary England. It was shown that the aggregate Intergenerational Earnings Elasticity (IEE) is low both in log-log and rank-rank specifications, likely downward biased by life-cycle effects. Additionally, using quantile regression of children's earnings on linear bins of parents' earnings illustrated that the IEEs vary across the joint distribution of parents' and children's earnings. In particular, IEEs are lower for higher levels of children's earnings but higher for higher levels of parents' earnings, a pattern which is not predicted by conventional models of human capital. To explain these results a model was constructed where parental earnings positively affects children's adulthood aspirations, and children invest in human capital to increase their earnings relative to their aspirations. This results in higher (lower) IEEs for children whose earnings are low (high) relative to their parents', consistent with the data.

There are several limitations and possible extensions of this analysis which naturally suggest avenues for future research. Firstly, this paper looked solely at earnings without directly investigating the mechanism through which they are increased. Such a model would predict differential responses to at-work incentives or education and training programs based on proximity to parental earnings – or even other potential determinants of aspirations – predictions which could be tested directly. Secondly, one could feasibly extend the model to a game-theoretical framework where earnings potential is a result of parents' decisions, who may also have an aspiration, and children can *then* choose to exert effort, labour supply or further human capital investment to make up the gap between their earnings and their aspirations.

However, taking these results at face value they represent an important finding. Two children who have similar ability, demographics and opportunity may have different life outcomes purely because one was born into a rich family and has higher aspirations whereas the other was born into a poor family and has lower aspirations. Policies which target aspirations could have high benefits for relatively low cost, while policies which reduce inequality may have additional benefits which have previously been underestimated. Future research could help to understand the magnitude of the effects of aspirations and the likely efficacy of different policies.

Appendices


A.1 Robustness Checks for Model Tests

This appendix uses different econometric specifications and measures of earnings to test the twin hypotheses predicted by the model: that IEEs decline as the quantile of children's earnings q rises, and that they rise as the parents' bin of earnings j increases. It uses an alternative earnings measure, has different cutoffs for the parentals' earnings splines, adds covariates, uses 3 splines instead of 2, and estimates a specification where the splines are joined. All specifications are consistent with the main results, though the lower spline $\hat{\theta}_{l,q}$ is imprecisely estimated in some cases.

Alternative Earnings Measure: Average parental earnings from 2004-2008 could be used as an alternative to the estimated lifetime earnings in the main tests of the aspirations model. Although this is not explicitly adjusted for life-cycle bias, it has several advantages. Firstly, it does not require imputation and therefore reduces the size of the standard errors. Secondly, this point in the life cycle may also be salient to children as it is when they make decisions about higher education, a huge factor in lifetime earnings. Thirdly, it is plausible that families are not able to estimate life cycle adjusted comparisons and so will use this information, which is more readily available.


Figures A1-A3 show that the main results are not affected by this choice. The downward trajectories of the quantile coefficients are arguably even more pronounced for observed earnings in 2007/08, and both $(\hat{\theta}_{l,75} - \hat{\theta}_{l,25})$ and $(\hat{\theta}_{h,75} - \hat{\theta}_{h,25})$ are significantly different at the 5% level in an interquantile regression. Although the coefficients are closer in magnitude than for the adjusted earnings measures, the standard errors are much smaller and the one-sided null hypothesis that $\hat{\theta}_{l,q} = \hat{\theta}_{h,q}$ is rejected at the 5% level for all quantiles except the 10th.

Figure A1: Robustness: Averaged Parental Earnings 2004-2008 ($\hat{\theta}_{l,q}$ & $\hat{\theta}_{h,q}$)




IEEs for both high and low earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The former is given by the solid grey line while the latter is given by the solid black line. Confidence intervals are not included on this graph for simplicity, but it is clear that the IEE for high earning parents is higher and formal tests reject the one-sided null hypothesis that the two are equal at every quantile. This graph uses averaged parental earnings over the sample rather than a life-cycle adjusted measure.

Figure A2: Robustness: Averaged Parental Earnings 2004-2008 ($\hat{\theta}_{h,q}$)



IEEs for high earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. This graph uses averaged parental earnings over the sample rather than a life-cycle adjusted measure.


Figure A3: Robustness: Averaged Parental Earnings 2004-2008 ($\hat{\theta}_{l,q}$)



IEEs for low earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. This graph uses averaged parental earnings over the sample rather than a life-cycle adjusted measure.

Different Cutoffs: Figure A4 shows the results when the 60th percentile of y^p is used as a cutoff and Figure A5 shows the results for the 40th percentile. For brevity both splines are shown on the same diagrams without confidence intervals, as in Figure 3.8. For the 60th percentile cutoff the results go through much as before: the one-sided test $H_0 : \theta_{l,75} = \theta_{l,25}$ can be rejected at 5% and the one-sided test $H_0 : \theta_{h,25} = \theta_{h,75}$ is rejected at 1%, while $H_0 : \theta_{h,q} = \theta_{l,q}$ is rejected in 8/9 quantiles.


Figure A4: Robustness: 60th Percentile Cutoff



IEEs for both high and low earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. Confidence intervals are not included on this graph for simplicity, but it is clear that the IEE for high earning parents is higher and formal tests reject the one-sided null hypothesis that the two are equal at every quantile. In this specification the upper spline captures values of parental earnings above the 60th percentile while the lower spline captures values below it.

Conversely, for the 40th percentile cutoff most of the tests do not go through. From inspection of Figure A5 it is clear that the two splines are close at every quantile and $H_0 : \theta_{h,q} = \theta_{l,q}$ is not rejected at any quantile. The one-sided test $H_0 : \theta_{l,25} = \theta_{l,75}$ cannot be rejected while the one-sided test $H_0 : \theta_{h,25} = \theta_{h,75}$ is rejected at only 10% confidence.


Figure A5: Robustness: 40th Percentile Cutoff



IEEs for high earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. Confidence intervals are not included on this graph for simplicity, but it is clear that the IEE for high earning parents is higher and formal tests reject the one-sided null hypothesis that the two are equal at every quantile. In this specification the upper spline captures values of parental earnings above the 40th percentile while the lower spline captures values below it.

Covariates: Figure A6 and Figure A7 show the results for both splines with the addition of the two types of covariates which were defined in the main analysis. The ‘exogenous’ covariates included in Figure A6 are race, gender and three ‘stated preference’ variables from the children: trust, risk and patience, all ranked on a scale of 1-10. These would not necessarily be expected to mediate the relationship with aspirations, which should hold for these groups. The ‘endogenous’ covariates included in Figure A7 are those which may serve as the margin on which children’s earnings are increased: education, health, and occupation. These therefore might be expected to mediate the aspirations relationship (although the available variables are not exhaustive, for example work-specific skills and training are not included).


Figure A6: Robustness: Exogenous Covariates



IEEs for both high and low earning parents at every quantile of children’s earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. Confidence intervals are not included on this graph for simplicity, but it is clear that the IEE for high earning parents is higher and formal tests reject the one-sided null hypothesis that the two are equal at every quantile. These regressions include ‘exogenous’ covariates race, gender, and three ‘stated preference’ variables from the children: trust, risk and patience, all ranked on a scale of 1-10.

In both figures the downward trajectory of the IEEs with quantiles is preserved even with covariates, and the upper spline is larger. Table A1 summarises the formal

Figure A7: Robustness: Endogenous Covariates



IEEs for both high and low earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. Confidence intervals are not included on this graph for simplicity, but it is clear that the IEE for high earning parents is higher and formal tests reject the one-sided null hypothesis that the two are equal at every quantile. These regressions include 'endogenous' covariates: education, health, and occupation.

hypothesis tests for both cases, with the original results from above, presented for comparison in column (1). In column (2) the addition of exogenous variables renders $\theta_{l,25}$ and $\theta_{l,75}$ statistically indistinguishable at conventional significance levels. Column (4) shows that the addition of endogenous covariates actually reduces both p-values.

The results from the null hypothesis of $H_0 : \theta_{l,q} = \theta_{h,q}$ against the one-sided alternative of $H_A : \theta_{l,q} < \theta_{h,q}$ is also reported in the first row of Table A1. The hypothesis is rejected at the 5% level for most quantiles with no covariates or exogenous covariates, but is not rejected at most quantiles with endogenous covariates. This is consistent with the idea that the endogenous covariates mediate the relationship between parents' and childrens' earnings.

Table A1: Hypothesis Tests with Covariates

	No Covariates	Exogenous Covariates	Endogenous Covariates
Hypothesis Test	(1)	(2)	(3)
$H_0 : \theta_{h,q} = \theta_{l,q}$	Rejected at 5% for 6/9 quantiles.	Rejected at 5% for 7/9 quantiles.	Rejected at 5% for 3/9 quantiles.
$H_A : \theta_{h,q} > \theta_{l,q}$			
$H_0 : \theta_{h,25} = \theta_{h,75}$	$p = 0.004$	$p = 0.055$	$p = 0.040$
$H_A : \theta_{h,25} > \theta_{h,75}$			
$H_0 : \theta_{l,25} = \theta_{l,75}$	$p = 0.015$	$p = 0.123$	$p = 0.086$
$H_A : \theta_{l,25} > \theta_{l,75}$			

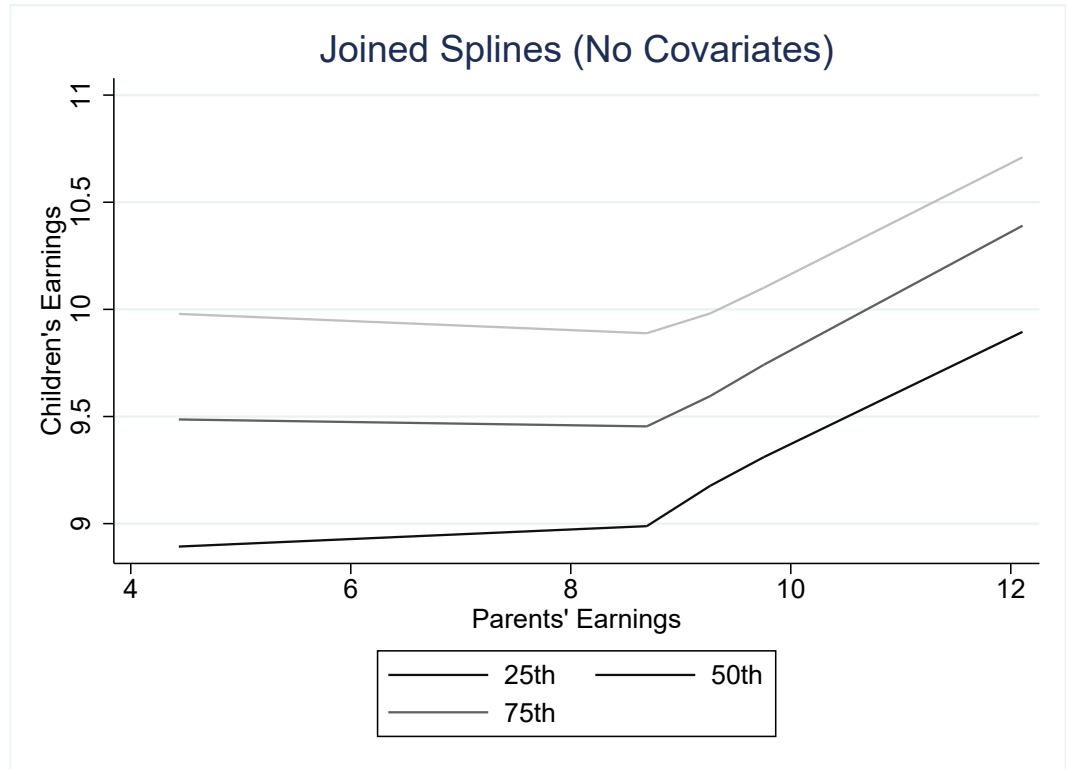
Results of the formal hypothesis tests of the two predictions of the aspirations model: that IEEs will decline as the quantile of children's earnings increases, and that the IEE for the upper spline will be higher than that for the lower spline. The first column shows the formal hypothesis tests while the next three columns show the results for no covariates, exogenous covariates and endogenous covariates.

Joined Splines An alternative specification ensures that the estimated splines will be joined at the knots so the function is continuous in y_i^p :

$$y_q^c = \theta_{0,q} + \theta_{1,q}y_i^p + \sum_{j=2}^J \theta_{j,q} \mathbb{1}(y_i^p > a_{j-1}) * (y_i^p - a_{j-1}) + \delta_q X_i + \varepsilon_{i,q} \quad (10)$$

The parameter $\theta_{1,q}$ is the elasticity of child's earnings at quantile q with respect to parents' earnings in bin 1, where $y^p \leq a_1$. The elasticity with respect to bin 2 is $(\theta_{1,q} + \theta_{2,q})$, where $a_1 < y^p \leq a_2$, the elasticity with respect to bin 3 is $(\theta_{1,q} + \theta_{2,q} + \theta_{3,q})$, and so on.

Figure A8: Linear Splines



IEEs estimated on 4 joined splines for the 25th, 50th, and 75th quantiles, with parents' earnings on the x-axis and children's earnings on the y-axis. Note that this graph is similar to the predictions of the aspirations model and the empirical results as summarised in Figure 3.5.

Figure A8 shows the results of regression (10) with $J = 4$ and for $q = 25, 50, 75$. Once more the IEEs are larger for higher values of parents' earnings, but smaller for higher values of children's earnings, though the latter is harder to see in the graphs than the tables. Table A2 shows the same results and it is clear that the slope is roughly zero

until the first inflection, after which the slope is roughly constant. Thus in this specification there seems to be a sudden change in the IEE rather than a gradual increase as in the case of unjoined splines. Table A2 shows the same results and with standard errors it is clear that only $\hat{\theta}_{2,q}$ is significant (for every q), further supporting the idea of a jump. It also declines as q rises, in line with the previous results and the aspirational model.

Table A2: IEEs Across the Distribution (Joined Splines)

	25th Quantile (1)	50th Quantile (2)	75th Quantile (3)
Log Parents' Earnings	0.0224 (0.0412)	-0.00205 (0.0460)	-0.0213 (0.0421)
Bin 1	0.304** (0.129)	0.245*** (0.0863)	0.182** (0.0796)
Bin 2	-0.0520 (0.201)	0.0577 (0.149)	0.0852 (0.135)
Bin 3	-0.0262 (0.176)	-0.0245 (0.153)	0.0111 (0.141)
Observations	5177	5177	5177

a. This table shows the elasticities from quantile regressions of children's earnings on bins of parents' earnings in the LSYPE as in regression (10). Each column shows a different quantile of children's earnings while each row shows a different bin of parents' earnings. In this estimation, bins are designed so that the splines join and children's earnings as a function of parents' earnings is continuous, with Bin 3 capturing the highest values of parents' earnings. The function is given by summing the coefficients, so that for parents' earnings in Bin 2 would be the coefficient on log parents' earnings + the coefficient on Bin 1 + the coefficient on Bin 2.


b. Robust standard errors in parentheses

c. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Tests with 3 Splines Two splines is the most straightforward way to test the implication of the aspirations model that $\theta_{h,q} > \theta_{l,q}$ for every q . However, that framework implies that the upper splines will be higher for any number of splines J so it is worth testing it for 3 splines (any more and there are issues with large standard errors due to low sample size).


Figures A11- A12 show the results of regression (3.4) with $J = 3$ and though it is clear that all 3 IEEs decline across the quantiles the standard errors are large and interquantile regression shows that $\theta_{j,25}$ and $\theta_{j,75}$ are not statistically distinguishable for every j . However, the hypothesis test $H_0 : \theta_{h,q} = \theta_{m,q} = \theta_{l,q}$ against the alternative $\theta_{j,q} \neq \theta_{j,q}$ for any $j = l, m, h$ is rejected at the 1% level at every quantile.

Figure A9: Robustness: 3 Splines ($\hat{\theta}_{l,q}$)




IEEs for lowest earning parents – defined as the bottom third – at every quantile of children’s earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. Formal hypothesis tests reject that the IEEs are equal between the 10th and 90th quantiles at the 10% level.

Figure A10: Robustness: 3 Splines ($\hat{\theta}_{m,q}$)




IEEs for middle earning parents – defined as the middle third – at every quantile of children’s earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. Formal hypothesis tests reject that the IEEs are equal between the 10th and 90th quantiles at the 10% level.

Figure A11: Robustness: 3 Splines ($\hat{\theta}_{h,q}$)



IEEs for highest earning parents – defined as the upper third – at every quantile of children’s earnings, from 10 to 90 in jumps of 10. The point estimate is given by the solid black line while the confidence intervals are the dashed grey lines. Formal hypothesis tests reject that the IEEs are equal between the 10th and 90th quantiles at the 10% level.

Figure A12: Robustness: Earnings Levels (All Splines)



IEEs for high, middle and low earning parents at every quantile of children's earnings, from 10 to 90 in jumps of 10. Confidence intervals are not included on this graph for simplicity, but formal tests reject the one-sided null hypothesis that the three are equal at every quantile at the 1% level.

Bibliography

- Mohammed Abdellaoui. Parameter-free elicitation of utility and probability weighting functions. *Management Science*, 46(11):1497–1512, 2000.
- Sule Alan, Thomas Crossley, and Hamish Low. Saving on a rainy day, borrowing for a rainy day. *Working Paper*, 2012.
- Muhktar Ali. Probability and utility estimates for racetrack bettors. *Journal of Political Economy*, 85:803–815, 1977.
- Jake Anders. The link between household income, university applications and university attendance. *Fiscal Studies*, 33(2):185–210, 2012.
- Eduardo Andrade, Fernando Veloso, Regina Madalozzo, and Sergio G. Ferreira. Do borrowing constraints decrease intergenerational mobility in brazil? a test using quantile regression. *Inspere Working Paper*, 2003.
- Michael Argyle. *The Psychology of Happiness*. Methuen & Co. Ltd, 1987.
- Orazio P. Attanasio and Guglielmo Weber. Consumption and saving: Models of intertemporal allocation and their implications for public policy. *Journal of Economic Literature*, 48:693–751, 2010.
- James Banks, Richard Blundell, and Sarah Tanner. Is there a retirement-savings puzzle? *The American Economic Review*, 4(88):769–788, 1998.
- Nicholas C. Barberis. Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27:173–196, 2013.
- Ian Bateman, Alistair Munro, Bruce Rhodes, Chris Starmer, and Robert Sugden. A test of the theory of reference-dependent preferences. *The Quarterly Journal of Economics*, 112(2):479–505, 1997.

- David Bayliss, Wendy Olsen, and Pierre Walthery. Well-being during the recession in the uk. *Working Paper*, 2014.
- Gary Becker and Nigel Tmes. An equilibrium theory of the distribution of income and intergenerational mobility. *Journal of Political Economy*, 87(6):1153–89, 1979.
- Gary Becker and Nigel Tmes. Human capital and the rise and fall of families. *Journal of Labor Economics*, 4(3):S1–S39, 1986.
- Chris Belfield, Claire Crawford, Ellen Greaves, Paul Gregg, and Lindsey Macmillan. Intergenerational income persistence within families. *IFS Working Paper W17/11*, 2017.
- Shlomo Benartzi and Richard H. Thaler. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110(1):73–92, 1995.
- Timothy Besley. Aspirations and the political economy of inequality. *Oxford Economic Papers*, 69(1):1–35, 2017.
- Sandra E. Black and Paul J. Devereux. *Handbook of Labor Economics*, chapter Recent Developments in Intergenerational Mobility. North Holland Press, Elsevier, 2011.
- Keith Blackburn and David Chivers. Fearing the worst: the importance of uncertainty for inequality. *Economic Theory*, 60(2):345–370, 2015.
- Jo Blanden and Lindsey Macmillan. Educational inequality, educational expansion and intergenerational mobility. *Journal of Social Policy*, 45(4):589–614, 2016.
- Teodora Boneva and Christopher Rauh. Parental beliefs about returns to educational investments—the later the better? *Journal of the European Economic Association*, 2018.
- Axel Borsch-Supan, Anette Reil-Held, and Reinhold Schnabel. *Life Cycle Savings and Public Policy*, chapter Household Saving in Germany, pages 57–99. Academic Press, 2003.
- David Bowman, Deborah Minehart, and Matthew Rabin. Loss aversion in a consumption-savings model. *Journal of Economic Behavior & Organization*, 38: 155–178, 1999.

- Rebecca R. Boyce, Thomas C. Brown, Gary H. McClelland, George L. Peterson, and William D. Schulze. An experimental examination of intrinsic values as a source of the wta-wtp disparity. *American Economic Review*, 82(5):1366–1373, 1992.
- Peter Brooks, Simon Peters, and Horst Zank. Risk behavior for gain, loss, and mixed prospects. *Theory and Decision*, 77(2):153–182, 2014.
- Martin Browning and Thomas F. Crossley. The lifecycle model of consumption and saving. *Journal of Economic Perspectives*, 15(3):3–22, 2001.
- Martin Browning and Annamaria Lusardi. Household saving: Micro theories and micro facts. *Journal of Economic Literature*, 34(4):1797–1855, 1996.
- Colin Camerer, Linda Babcock, George Loewenstein, and Richard Thaler. Labor supply of new york city cabdrivers: One day at a time. *The Quarterly Journal of Economics*, 1997.
- Colin F. Camerer. *Prospect Theory In The Wild: Evidence From The Field*. 2000.
- John Campbell and Angus Deaton. Why is consumption so smooth? *Review of Economic Studies*, 56:357–374, 1989.
- Centre for Longitudinal Studies. Next steps: Sweeps 1-8, 2004-2016. [data collection]. 14th edition, 2018. URL <http://doi.org/10.5255/UKDA-SN-5545-6>.
- Raj Chetty and Adam Szeidl. Consumption commitments: A foundation for reference-dependent preferences and habit formation. *Working Paper*, 2009.
- Raj Chetty, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129(4):1553–1623, 2014.
- Andrew E. Clark. Are wages habit-forming? evidence from micro data. *Journal of Economic Behavior & Organization*, 3(2):179–200, 1999.
- Andrew E. Clark and Andrew J. Oswald. Satisfaction and comparison income. *Journal of Public Economics*, pages 359–381, 1995.
- Andrew E. Clark, Paul Frijters, and Michael Shields. Relative income, happiness and utility: An explanation for the easterlin paradox and other puzzles. *Journal of Economic Literature*, pages 95–144, 2008.

- Ofer Cohen. Consumption smoothing with state-dependent loss-aversion preferences. *Mimeo*, 2015.
- George M. Constantinides. Habit formation: A resolution of the equity premium puzzle. *The Journal of Political Economy*, 9(3):519–543, 1990.
- Vincent P. Crawford and Juanjuan Meng. New york city cab drivers’ labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review*, 101:1912–1932, 2011.
- Patricio S. Dalton, Sayantan Ghosal, and Anandi Mani Anandi Mani, Sayantan Ghosal. Poverty and aspirations failure. *Economic Journal*, 126(590):165–188, 2016.
- Jan-Emmanuel De Neve, George W. Ward, Femke De Keulenaer, Bert Van Landeghem, Georgios Kavetsos, and Michael I. Norton. The asymmetric experience of positive and negative economic growth: Global evidence using subjective well-being data. *Review of Economics and Statistics*, 100(2):362–375, 2018.
- Lorraine Dearden, Stephen Machin, and Howard Reed. Intergenerational mobility in britain. *Economic Journal*, 107(440):47–66, 1997.
- Stefano Dellavigna, Attila Lindner, Balazs Reizer, and Johannes F. Schmieder. Reference-dependent job search: Evidence from hungary. *The Quarterly Journal of Economics*, 4:1969–2018, 2017.
- Edward Diener. Subjective well-being. *Psychological Bulletin*, pages 542–575, 1984.
- James Duesenberry. *Income, Saving and the Theory of Consumer Behaviour*. Harvard University Press, 1949.
- Richard A. Easterlin, Laura Angelescu McVey, Malgorzata Switek, Onnicha Sawangfa, and Jacqueline Smith Zweig. The happiness–income paradox revisited. *PNAS*, page 22463–22468, 2010.
- William Easterlin. *Does Economic Growth Improve the Human Lot? Some Empirical Evidence*. 1974.
- Eric R. Eide and Mark H. Showalter. Factors affecting the transmission of earnings across generations: A quantile regression approach. *Journal of Human Resources*, 34(2):253–267, 1998.

- Henry S. Farber. Is tomorrow another day? the labor supply of new york city cab-drivers. *Journal of Political Economy*, 113(1):46–82, 2005.
- Henry S. Farber. Reference-dependent preferences and labor supply: The case of new york city taxi drivers. *The American Economic Review*, 98(3):1069–1082, 2008.
- Amy Finkelstein, Erzo F. P. Luttmer, and Matthew J. Notowidigdo. What good is wealth without health? the effect of health on the marginal utility of consumption. *Journal of the European Economic Association*, 11(S1):221–258, 2012.
- Dieter Von Fintel. Earnings bracket obstacles in household surveys – how sharp are the tools in the shed? *Stellenbosch Economic Working Papers*, 2006.
- Bruno Frey and Alois Stutzer. The economics of happiness. *World Economics*, pages 25–41, 2002.
- Bruno Frey and Alois Stutzer. Recent advances in the economics of individual subjective well-being. *Social Research*, pages 679–714, 2010.
- B.S. Frey and A Stutzer. *Economics of Happiness*. Springer International Publishing, 2018.
- David Gal and Derek D. Rucker Derek D. Rucker. The loss of loss aversion: Will it loom larger than its gain? *Journal of Consumer Psychology*, 28(5):497–516, 2018.
- Garance Genicot and Debraj Ray. Aspirations and inequality. *Econometrica*, 85(2): 489–519, 2017.
- Richard Gonzalez. On the shape of the probability weighting function. *Cognitive Psychology*, 38:129–166, 1999.
- Nathan D. Grawe. Reconsidering the use of nonlinearities in intergenerational earnings mobility as a test for credit constraints. *The Journal of Human Resources*, 39(3): 813–827, 2004.
- Nathan D. Grawe. Lifecycle bias in estimates of intergenerational earnings persistence. *Labour Economics*, 13(5):551–570, 2006.
- Wim Groot. Adaptation and scale of reference bias in self-assessments of quality of life. *Journal of Health Economics*, 19(3):403–420, 2000.

- Alessandra Guariglia. Understanding saving behaviour in the uk: Evidence from the bhps. *Unpublished Manuscript*, 1998.
- Alessandra Guariglia and Mariacristina Rossi. Consumption, habit formation, and precautionary saving: evidence from the british household panel survey. *Oxford University Press*, 1(54):1–19, 2002.
- Song Han and Casey B. Mulligan. Human capital, heterogeneity and estimated degrees of intergenerational mobility. *The Economic Journal*, 111:207–243, 2001.
- Bo E. Honoré. Trimmed lad and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica*, 60(3):533–565, 1992.
- John Horowitz, John List, and K. E. McConnell. A test of diminishing marginal value. *Economica*, 74:650–663, 2007.
- Institute for Social and Economic Research. British household panel survey: Waves 1-18, 1991-2009. [data collection]. 8th edition., 2018. URL <http://doi.org/10.5255/UKDA-SN-5151-2>.
- Institute for Social and Economic Research, University of Essex, NatCen Social Research, and Kantar Public. Understanding society: Waves 1-9, 2009-2018 and harmonised bhps: Waves 1-18, 1991-2009. [data collection]. 12th edition., 2019. URL <http://doi.org/10.5255/UKDA-SN-6614-13>.
- Tullio Jappelli and Luigi Pistaferri. The consumption response to income changes. *Annual Review of Economics*, (2):479–506, 2010.
- Robert Jensen. The (perceived) returns to education and the demand for schooling. *Quarterly Journal of Economics*, 125(1):515–548, 2010.
- Andrew M. Jones and John Wildman. Health, income and relative deprivation: Evidence from the bhps. *Journal of Health Economics*, pages 308–324, 2008.
- Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, pages 263–291, 1979.
- Daniel Kahneman, Peter P. Wakker, and Rakesh Sarin. Back to bentham? explorations of experienced utility. *The Quarterly Journal of Economics*, pages 375–406, 1997.

- John Maynard Keynes. *The General Theory of Employment, Interest and Money*. Palgrave Macmillan, 1936.
- Miles Kimball and Robert Willis. Utility and happiness. *Presented at an IFS Conference*, 2006.
- Jack L. Knetsch and J. A. Sinden. Willingness to pay and compensation demanded: Experimental evidence of an unexpected disparity in measures of value. *Quarterly Journal of Economics*, 99(3):507–521, 1984.
- Gundi Knies. Understanding society –uk household longitudinal study: Wave 1-4, 2009-2013, user manual. *Colchester: University of Essex.*, 2014.
- Botond Koszegi and Matthew Rabin. A model of reference-dependent preferences. *The Quarterly Journal Of Economics*, 121(4):1133–1165, 2006.
- Botond Koszegi and Matthew Rabin. Reference-dependent risk attitudes. *American Economic Review*, 97(4):1047–1073, 2007.
- Botond Koszegi and Matthew Rabin. Reference-dependent consumption plans. *American Economic Review*, 99(3):909–936, 2009.
- Rachid Laajaj. Endogenous time horizon and behavioral poverty trap: Theory and evidence from mozambique. *Journal of Development Economics*, 127:187–208, 2017.
- John A. List. Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, 72(2):615–625, 2004.
- Yusufcan Masatlioglu and Collin Raymond. A behavioral analysis of stochastic reference dependence. *American Economic Review*, 9(9):2760–82, 2016.
- Alexander Michaelides. Buffer stock saving and habit formation. *Unpublished*, 2002.
- Franco Modigliani and Richard Brumberg. Utility analysis and the consumption function: An interpretation of cross-section data. *The Collected Papers of Franco Modigliani*, 6:388–436, 1955.
- Andrew J. Oswald. Happiness and economic performance. *Economic Journal*, pages 1815–1831, 1997.

- Andrew J. Oswald and Nattavudh Powdthavee. Does happiness adapt? a longitudinal study of disability with implications for economists and judges. *Journal of Public Economics*, pages 1061–1077, 2008.
- Michaela Pagel. Expectations-based reference-dependent life-cycle consumption. *Review of Economic Studies*, 84(2):885–934, 2016.
- Michaela Pagel. A news-utility theory for inattention and delegation in portfolio choice. *Econometrica*, 2018.
- D. J. Pevalin. Multiple applications of the ghq-12 in a general population sample: an investigation of long-term retest effects. *Soc Psychiatry Psychiatr Epidemiol*, pages 508–12, 2000.
- Robert A. Pollak. Habit formation and dynamic demand functions. *Journal of Political Economy*, 78(4):745–763, 1970.
- John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323–343, 1982.
- Matthew Rabin. Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5):1281–1292, 2000.
- G. Rajabi and S. Hashemi Sheykhshabani. Factor structure of the 12-item general health questionnaire. *Journal of Education & Psychology*, pages 81–94, 2009.
- Ulrich Schmidt, Chris Starmer, and Robert Sugden. Third-generation prospect theory. *Journal of Risk and Uncertainty*, 36:203–223, 2008.
- Arthur Seibold. Reference dependence in retirement behavior: Evidence from german pension discontinuities. *Working Paper, University of Mannheim*, 2017.
- John Shea. Does parents’ money matter? *Journal of Public Economics*, 77:155–184, 2000.
- Andrei Shleifer. Psychologists at the gate: A review of daniel kahneman’s thinking, fast and slow. *Journal of Economic Literature*, 50(4):1080–1091, 2012.
- Chris Starmer. Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *Journal Of Economic Literature*, 38(2):332–382, 2000.

- Betsey Stevenson and Justin Wolfers. Economic growth and subjective well-being: Reassessing the easterlin paradox. *Brookings Papers on Economic Activity*, pages 1–87, 2008.
- Marcia Freed Taylor, John Brice, Nick Buck, and Elaine Prentice-Lane. British household panel survey user manual volume a: Introduction, technical report and appendices. *Colchester: University of Essex.*, 2010.
- James Tobin. Estimation of relationships for limited dependent variables. *Econometrica*, 26(1):24–36, 1958.
- Peter Warr. Reported behaviour changes after job loss. *British Journal of Social Psychology*, pages 271–275, 1984.
- Eldad Yechiam. Acceptable losses: the debatable origins of loss aversion. *Psychological Research*, page 1–13, 2018.