University of Oxford

Doctoral Thesis

The long-run effects of public policy on human capital and savings

Ph.D. candidate: Nick Ridpath

Supervisors: Prof. Hamish Low
Dr. Sarah Clifford

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy in Economics

September 2024

Abstract

This thesis examines the long-run effects of a range of public policies on household decision-making, and their consequences for saving and human capital later in life. The first chapter looks at the effect of childcare subsidies on the decisions and outcomes of mothers, using a dynamic life-cycle model to estimate the labour supply and savings response to reduced childcare costs. This paper finds that a hypothetical universal childcare policy would significantly increased hours worked for mothers who were already in work, with a small but positive effect on earnings in the long run. The second chapter uses a similar model to estimate the savings response to changes to adult social care, examining how household consumption responds to changes to the level of insurance against social care costs individuals receive. The work outlines how greater social insurance would not trigger a larger saving response, but would be welfare-improving. The third chapter uses a difference-in-differences strategy to estimate the effect of a conditional cash transfer for young people staying in education on long-run outcomes, finding that the additional incentive to stay in school did not lead to increased wages later in life.

Word Count
This thesis consists of approximately 31,000 words.

Acknowledgements

This thesis would not have been possible without the support I have received over the last four years. I would like to thank Hamish Low and Sarah Clifford for their teaching and supervision, which gave me the skills and helped me develop the ideas to put this piece of work together. Teaching and supervision from Miguel Ballester during my MPhil, and Bassel Tarbush, Michalis Rousakis and Sergi Pardos-Prado during my undergraduate degree, helped guide me and inspire me to this point.

In addition to intellectual support, I have benefited greatly from the financial support I received throughout my DPhil. Without the support of University College through the Oxford-Chellgren Graduate Scholarship, the Department of Economics, and the Institute for Fiscal Studies, this thesis would not exist. I and my co-authors at the IFS are also grateful for funding from the Nuffield Foundation (grant EDO/FR-000023448). The Nuffield Foundation is an independent charitable trust with a mission to advance social wellbeing. It funds research that informs social policy, primarily in Education, Welfare and Justice. The Nuffield Foundation is the founder and co-founder of the Nuffield Council on Bioethics, the Ada Lovelace Institute and the Nuffield Family Justice Observatory. The Nuffield Foundation has funded this project, but the views expressed are those of the authors and not necessarily the Foundation. Website: www.nuffieldfoundation.org. X (Twitter): @NuffieldFound. Co-funding from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy (ES/T014334/1) is gratefully acknowledged.

I do not believe you can write learn how to finish a DPhil on your own, and I am extraordinarily grateful to the people around me during the DPhil, who taught me what they had already learned, or learned alongside me. Shihang, Luke, Hannah, Binta, Sam, Alex, and Stan and many others made the process so much easier, and made long days in the Economics Department much more fun. Thank you also to Gwion, Fran, Rebecca, and Matt, often the only people I would see each day while writing a thesis in the pandemic, who helped keep me sane. At the IFS, Ben, Jack, Carmen, Bee, and Sarah, among others, provided the support I needed to get the thesis over the line.

My family have been an enormous help throughout this process. In particular, thank you to my parents, who first fostered my interest in economics, and who supported me and taught me so much of what I know.

And finally, thank you to Niamh, who was by my side this whole time. A DPhil is often cause for a lot of stress, nerves, and complaining, and she not only put up with it, but gave me the confidence to believe in what I was doing. I do not know how I would have managed without her.

Contents

List of Figures **ix**

List of Tables **xi**

Introductory Remarks **1**

1 Childcare subsidies and the child penalty: a life-cycle analysis **3**

1.1 Introduction . . . 3
1.2 Life-Cycle Model . . . 6
1.2.1 Problem of the individual . . . 6
1.2.2 Childcare . . . 7
1.2.3 Maximisation Problem . . . 8
1.2.4 Wage Processes . . . 9
1.2.5 Exogenous Variables . . . 10
1.2.6 Solution . . . 11
1.3 Data . . . 11
1.4 Estimation . . . 13
1.4.1 Externally Set Parameters . . . 13
1.4.2 Directly Estimated Parameters . . . 13
1.4.3 Estimation by Simulated Method of Moments . . . 15
1.4.4 Fit of the Model . . . 19
1.5 Counterfactual Results . . . 21
1.5.1 Lifetime Labour Market Outcomes . . . 22
1.5.2 Child Penalty . . . 25
1.5.3 Welfare Analysis . . . 28
1.6 Conclusion . . . 29

2 Adult social care funding and its interaction with saving for retirement **31**

2.1 Introduction . . . 31
2.2 Institutional Setting . . . 34
2.3 Life-Cycle Model . . . 34
2.3.1 Problem of the Individual . . . 35
2.3.2 Health Process and Death . . . 36
2.3.3 Social Care . . . 37
2.3.4 Maximisation Problem . . . 39
2.3.5 Wages . . . 42
2.3.6 Solution . . . 42
2.4 Data . . . 43
2.4.1 Initial Conditions . . . 43
2.5 Estimation . . . 44
2.5.1 Externally Set Parameters . . . 44
2.5.2 Directly Estimated Parameters . . . 44
2.5.3 Estimation by Simulated Method of Moments . . . 49
2.5.4 Retirement Moments . . . 51
2.5.5 Housing Moments . . . 52
2.5.6 Bequest and Care Moments . . . 52
2.6 Counterfactuals . . . 53
2.6.1 Comparison of Social Care Funding by Generosity . . . 54
2.6.2 Comparison by Nature of Housing Test . . . 57
2.6.3 Welfare Analysis . . . 59
2.7 Conclusion . . . 61

3 The long-run effect of the Education Maintenance Allowance 63
3.1 Introduction . . . 63
3.2 Institutional Background . . . 66
3.2.1 Post-16 Education in England . . . 68
3.3 Data . . . 69
3.4 Methodology . . . 72
3.5 Results . . . 74
3.5.1 Effect of the EMA on education participation . . . 74
3.5.2 Effect on short-run labour supply . . . 78
3.5.3 Effect of the EMA on qualifications . . . 79
3.5.4 Effect of the EMA on long-run earnings . . . 80
3.6 Conclusion . . . 82

Concluding Remarks **85**

Appendix A Appendix to: Childcare subsidies and the child penalty: a life-cycle analysis **89**
A.1 Solution of the Model . . . 89
A.1.1 Use of the Endogenous Grid Method . . . 89
A.1.2 Details of Solution and Simulation . . . 90
A.2 Estimation . . . 91
A.2.1 Estimation of Male Income Process . . . 91
A.2.2 Marriage Parameters . . . 92
A.2.3 Fertility Probabilities . . . 92
A.3 Model Fit . . . 93
A.3.1 List of Moments . . . 93
A.3.2 Relevance of Moments . . . 95

Appendix B Appendix to: Adult social care funding and its interaction with saving for retirement **101**
B.1 Solution of the Model . . . 101
B.1.1 Use of the Endogenous Grid Method . . . 101
B.1.2 Details of Solution and Simulation . . . 102
B.2 List of instruments for self-reported health . . . 102
B.3 Model Fit . . . 103
B.3.1 List of Moments . . . 103
B.3.2 Relevance of Moments . . . 105

Appendix C Appendix to: Adult social care funding and its interaction with saving for retirement **107**
C.1 Standard difference-in-differences results . . . 107

References **111**

List of Figures

1.1 Ratio of Mothers' Labour Market Statistics relative to Non-Mothers of the same age in the model and the data 22
1.2 Labour Supply and Income Over Time 23
1.3 Employment and Income Over Time Amongst Those Who Have to Pay for Childcare 25
1.4 Ratio of Mothers' Labour Market Statistics relative to Non-Mothers of the same age 26

2.1 Savings Rates Over Time 55
2.2 Log Assets Over Time 57
2.3 Home Ownership Over Time 58
2.4 Portfolio Allocation Over Time 59

3.1 Education Maintenance Allowance Rollout 67
3.2 Effect on Proportion Earning More Than A Certain Level At Different Ages 83

List of Tables

1.1 Distribution of Family Makeups, Age 20-40 12
1.2 Women's Employment by Family Makeup, Age 20-65 12
1.3 Utility Cost of Work Parameters by Number of Children, Age of Children and Marital Status 16
1.4 Income Process Parameters 18
1.5 Human Capital Parameters 18
1.6 Paid Childcare Probability 19

2.1 Probability of Needing Any Care 46
2.2 Probability of Needing Residential Care 47
2.3 Number of Hours of Care Needed 48
2.4 Income Moments 50
2.5 Income Parameters 50
2.6 Retirement Moments 51
2.7 Retirement Parameters 51
2.8 Housing Moments 52
2.9 Housing Parameters 52
2.10 Other Moments 53
2.11 Other Parameters 53
2.12 Average Retirement Age, by Generosity of Policy 56
2.13 Tax Required to Make Each Policy Revenue Neutral 60
2.14 Tax Required to Make Each Policy Revenue Neutral 60

3.1 Background Characteristics of Students on Free School Meals, Born 1985-1995 71
3.2 Age 16 Outcomes of Students on Free School Meals, Born 1985-1995 . . . 71
3.3 Impact of the EMA on education participation aged 16-17 (FSM students) . 74
3.4 Impact of the EMA on education participation aged 17-18 (FSM students) . 76
3.5 Impact of the EMA on education participation by gender (FSM students) . . 77

3.6 Impact of the EMA on education participation by prior academic performance (FSM students) . . . 77
3.7 Impact of the EMA on education participation by special educational needs (FSM students) . . . 78
3.8 Impact of the EMA on work at age 17 (FSM students) . . . 79
3.9 Impact of the EMA on qualifications (FSM students) . . . 81
3.10 Impact of the EMA on earnings (FSM students) . . . 82

A.1 Male Income Process Parameters . . . 91
A.2 Marriage Probability Parameters . . . 92
A.3 Probability of Childbirth . . . 93
A.4 Employment Coefficient by Family Demographics . . . 94
A.5 Remaining Employment Moments . . . 95
A.6 Human Capital Regressions . . . 95
A.7 Income Statistics . . . 96
A.8 Other Moments . . . 99

B.1 Link between different diseases and self-reported health . . . 103
B.2 Moments in Data and Model . . . 104

C.1 Impact of the EMA on education participation aged 16-17 (FSM students) . . . 107
C.2 Impact of the EMA on education participation aged 17-18 (FSM students) . . . 108
C.3 Impact of the EMA on education participation by gender (FSM students) . . . 108
C.4 Impact of the EMA on education participation by prior academic performance (FSM students) . . . 109
C.5 Impact of the EMA on education participation by special educational needs (FSM students) . . . 109
C.6 Impact of the EMA on work at age 17 (FSM students) . . . 109
C.7 Impact of the EMA on qualifications (FSM students) . . . 110
C.8 Impact of the EMA on earnings (FSM students) . . . 110

Introductory Remarks

The long-run effect of policy is often key to the impact of public policy, but the evaluation of impacts of a policy over the course of an individual's life is very challenging. Data limitations and a lack of clean identification can hinder the estimation of the intended effects of some policies, whether it be improving lifetime earnings, or insuring households in response to risks far in the future. Using a combination of structural methods and administrative datasets, this thesis aims to overcome these challenges, answering three different questions over three chapters around the effect of policies with intended long-run effects.

The first chapter focuses on the effect of childcare subsidies on the productivity and earnings of mothers throughout their working life. While childcare subsidies can have immediate benefits by reducing the disincentive to work for mothers who would choose to work absent childcare costs, we might expect there to be significant benefits in the long run - remaining attached to the labour force might be expected to increase human capital over the next several decades. The second chapter, focused on adult social care reforms, considers an inverted question of the role of long term thinking in policy - how does a policy that insures households against risks far in the future affect their decisions now? Specifically, the chapter looks at how household adjust how much they save, and in what form they save, when they can expect a reduced risk of having to spend large sums on social care in their old age. The third chapter, meanwhile, looks at a policy specifically intended to have long-run effects. Cash transfers to young people conditional on education participation are often introduced with the aim of boosting skills and productivity in the long run, and so it is important to understand what they do for productivity as young people reach adulthood. This chapter studies one such policy, the Education Maintenance Allowance, using longitudinal administrative data.

The first chapter demonstrates the importance of the consideration of long-run impacts when thinking about childcare policy. Few settings have data required to evaluate the effect of childcare subsidies over a long timescale. To get around this, I construct a life-cycle model of maternal labour supply and household saving, to estimate the long-run returns to reduced

childcare costs. It finds that human capital accumulation through labour market experience is an important channel, and so that policies like childcare subsidies that can increase labour supply in the short run can have long-run benefits for wages, as well as allowing mothers to earn more in the short run through increased hours in work. I find that even when taxes are levied to pay for this policy, it can be welfare-improving, precisely because of these long-run benefits - the strong disincentive to work for many new mothers stemming from high childcare costs can hinder their earnings for a long time, and removing that disincentive can significantly increase lifetime earnings.

The second chapter highlights the necessity of evaluating social insurance policies from an ex-ante perspective, and the difficulty of doing so when the insurance is against risks occuring late in life. Building a life-cycle model of saving for retirement in the context of different policies for social care subsidies, I evaluate how different levels of protection against very high social care costs affect households decisions and ex-ante welfare from the age of 50 onwards. The results show that additional support with social care from the government would be welfare-improving, but would not significantly change saving or retirement decisions of households in advance. However, reductions from current levels of social care support would drastically affect household saving: if government support for social care costs only went to those with no assets, the government would receive, we should expect to see very large increases in household saving, and all families retiring much later, demonstrating clearly the role social insurance against such risks can play over the long term.

Instead of using structural modelling, the third chapter uses large administrative data to get around the difficulty of estimating the long-run effects of policy. The Education Maintenance Allowance in the UK, as well as other cash transfers conditional on education like it worldwide, target long-run benefits for affected young people by boosting their skills for later life. Using data that links education choices and outcomes to earnings up to age 30, this chapter is able to estimate whether the short run effects of such subsidies necessarily translate to long-run effects. While the paper does find an increase in education participation when there is additional financial incentive to take part, it does not translate to additional earnings in the long run. This is likely because the courses which saw an increase in takeup due to the EMA were generally of low quality, and did not increase earnings. This emphasises the importance of being careful in assuming long-run benefits from positive short term indicators, and the necessity of evaluating policy from a long-term perspective where possible.

1 Childcare subsidies and the child penalty: a life-cycle analysis

Abstract

The "child penalty" of motherhood represents the reduction in women's wages upon having children. This paper considers the question of how much of this reduction would be affected by a universal childcare subsidy. I do this by estimating a rich life-cycle model of women's labour supply and savings choices, using data from Understanding Society. This considers the utility cost of work, the human capital accumulation process, and access to informal forms of childcare. Using this model, I estimate the effect of universal childcare on this child penalty, as well as on women's wages over the long-run. I find that the cost to women's earnings from having children is significantly reduced with a universal childcare policy, and that such a policy leads to a small increase in women's wages later in life.

1.1 Introduction

The role of children in affecting female labour market outcomes is an important one. Studies have shown that the "child penalty" in earnings for women having children is large and takes the form of a reduction in wages and a reduction in labour market participation, at both the intensive and extensive margin over the long term (Kleven et al., 2019). One plausible driver of this child penalty is childcare costs, which add an additional cost to labour market participation for parents of young children. In this paper, I study how policies reducing

childcare costs would affect this child penalty. I also examine how it would affect female labour market participation and long-run wages in the aggregate, across both mothers and non-mothers.

To do this, I build a life-cycle model of women's consumption and labour market decisions. This allows us to understand the determinants of women's labour supply decisions before, during and after childcare. The use of a dynamic model allows me to illuminate the role labour supply choice plays in human capital accumulation, the key mechanism behind the interaction between employment choices in the early years of motherhood and wages in the long run. The use of a structural model allows me to estimate three key sets of unobservable parameters using simulated method of moments: the non-pecuniary costs of work dependent on children and marital status, the parameters governing human capital accumulation, and the proportion of women who need to use formal channels to acquire childcare.

Estimating these parameters allows me to examine how counterfactual childcare subsidies would impact female labour market participation. I simulate these counterfactuals and estimate how they would affect employment during the early years of motherhood. By estimating the human capital accumulation process, and by studying labour market outcomes across women's working life, I then consider how long-run wages are affected by the subsidy. I measure this both in terms of the effect on the gap between mothers and non-mothers, and in terms of the aggregate impact on wages across all women.

I use survey data from Understanding Society to estimate the labour market and saving choices of non-graduate women in the United Kingdom, as well as how those choices interact with future labour market outcomes. I do this by estimating a model of choices over the course of women's working life, using simulated method of moments to match key facts in the data to equivalent elements of the simulated model. This allows me to estimate important unobservable parameters. Using these unobservable parameters, I then run the same simulated model while including various childcare subsidy scenarios, and observe the change in labour market outcomes, both in terms of aggregate outcomes across all women over the life cycle and in terms of the child penalty.

The results show that the government completely subsidising the cheapest types of market childcare would have large effects on life-cycle labour market outcomes. I find that a full childcare subsidy funded by income tax can increase mean employment for non-graduate women by as much as 7 percentage points at certain points of the life cycle. I also find that it can have a positive impact on long-run mean wages. The results are driven by the subset of the population that does not have access to any other type of childcare, such as from grandparents. This is because only their labour market participation would be affected by the cost of market childcare.

The results also show significant impact on the child penalty, which I evaluate by comparing the labour market outcomes in a comparison between mothers and non-mothers of the same age. I find that the penalty in total labour earnings from having children, which is greater than 30% for much of motherhood and remains above 20% 20 years after becoming a mother, is reduced by between a quarter and a half in every year from becoming a mother to the first child turning 20. This is driven by increases on all three margins of total earnings: employment, hours worked and wages.

Various strands of the life-cycle literature have modelled women's labour market outcomes over the life cycle. Adda et al. (2017) look at how having children affects human capital accumulation through the lens of occupation choice. Eckstein et al. (2019) use a life cycle model to examine how childbirth decisions have affected changing wages in the post-war period. Similarly, Attanasio et al. (2008) look at how childcare cost changes in the second half of the 20th century changed the gender pay gap. Meanwhile, Blundell et al. (2016) use labour market participation choices to look at the role of Working Tax Credit reforms in the United Kingdom.

Meanwhile, other life-cycle work has focused on childcare policy specifically, including Bernal (2008), Chan and Liu (2018), and Boca et al. (2014). While much of this work does touch on labour supply responses to childcare policy changes, it is generally in the short run, with the long-run effects focused more closely on the effects on child development.

More recently, a significant literature has developed estimating the size of the child penalty. Much of the work has come from event study frameworks (Casarico and Lattanzio, 2023; Kleven et al., 2019, 2023), while other evidence comes from IV-based strategies exploiting unexpected fertility shocks (Bensnes et al., 2023; Bögl et al., 2024), generally finding smaller effects. Other work on the child penalty has explored the role of policies such as parental leave (Adams et al., 2024) or work search requirements (Kuka and Shenhav, 2024), or explored differential effects by education (Goldin et al., 2024).

There have also been examinations of natural experiments arising from policy changes to childcare. Significant reduced-form work has shown that policy changes reducing childcare costs and increasing availability have had significant positive impact on mothers' labour market participation in England (Brewer et al., 2022), Spain (Nollenberger and Rodríguez-Planas, 2015) and Italy (Carta and Rizzica, 2018). This paper presents similar results, but I go on to show how this effect would stretch out by pushing up mothers' wages in the long run.

The rest of the paper is as follows. Section 2 outlines the life-cycle model, showing the structure of utility, wages and human capital accumulation that drives decision making in the model. Section 3 outlines the data used to estimate the model. Section 4 outlines the

estimation process. This includes details of the estimation of exogenous parameters, and an outline of the method and moments used to match the unobservable parameters in the model. Section 4 also includes the results of this estimation, outlaying the estimates of these unobservable parameters. Section 5 demonstrates the results from the counterfactual analysis, through measurements of employment, wages, the child penalty, and a welfare analysis. Section 6 concludes.

1.2 Life-Cycle Model

The model in this paper follows the saving and labour supply decisions of adult women over the course of their working life. At the beginning of each period, women observe a series of shocks, including productivity, fertility, and marital shocks. Upon observing these shocks, they make decisions how much to work, and how much money to save and consume. If they have children, their labour supply decision has consequences for their use of childcare, and for how much they pay for childcare. The model also accounts for human capital accumulation, meaning that labour supply decisions affect future productivity. The aim of the model is to examine how women adjust their labour supply decisions when the cost of childcare changes, and to think about the implications of this for women's long-term earnings.

The period modelled is from age 20 to 65. Each period t in the model represents their consumption and saving decisions at age t. At age 65, women retire and, after that point, can no longer make any decision to work, only a consumption decision.[1] After retirement, women live for ten more years from their accumulated savings.

1.2.1 Problem of the individual

At age t, women make decisions on how much to work ($P_{i,t} \in \{0, 0.5, 1\}$), and how much to consume ($c_{i,t}$). The different values of P represent proportions of a 40-hour week spent working. Marriage and childbirth are determined by exogenous shocks. The decisions made by the individual take place after observing marital status, childbirth, and productivity shocks. These decisions are taken based on the individual's preferences, which are represented by $u(P_{i,t}, c_{i,t})$. Additional state variables savings ($A_{i,t,}$) and human capital ($h_{i,t}$) are dependent on the consumption and labour supply decisions made. If the woman i is single, then her

[1]There is no possibility of change of marital status after 65, so women either stay married or stay single until death. There are also no bequests in the model, so women spend all their savings before they die.

budget constraint is as follows:

$$\frac{A_{i,t+1}}{1+r} = A_{i,t} - c_{i,t} + (w_{i,t} - Paid_i * CC_{i,t})P_{i,t} + B_{i,t} \tag{1.1}$$

$$A_{i,t+1} \geq 0 \tag{1.2}$$

where t is her age. The wage rate is determined by a persistent stochastic process, and by levels of human capital, as outlined below. $B_{i,t}$ is the level of benefits received, which is given by a formula that follows the UK benefits system in 2011. It is a function of number of children, age of children, employment status, income, and assets. $CC_{i,t}$ represents childcare costs, and is also dependent on the number and age of children. $Paid_{i,t}$ is a binary variable representing whether a woman received informal support with childcare.

Meanwhile, if the woman is married, then the decisions to be made are how much the woman works and how much the household consumes. The budget constraint for the married household is therefore as follows:

$$\frac{A_{i,t+1}}{1+r} = A_{i,t} - c_{i,t} + (w_{i,t} - Paid_i * CC_{i,t})P + w^M_{i,t} + B_{i,t} \tag{1.3}$$

$$A_{i,t+1} \geq 0 \tag{1.4}$$

where $w^M_{i,t}$ is the income of her spouse. I follow much of the female labour supply life-cycle literature (Adda et al., 2017; Blundell et al., 2016; Low et al., 2018) and set male labour supply to be exogenous, as well as assuming that men always have some form of income, whether through labour or benefits. In doing so, I effectively take the gender dynamics that may drive the gender wage gap and child penalty as a given, and consider policies that might alter these outcomes under these norms.

1.2.2 Childcare

Childcare is represented by $CC_{i,t}$ in the model. This is paid for if the woman works, and if they have no alternate option for childcare. The binary variable $Paid_i$ represents whether an individual has access to an alternative option for childcare, such as a relative or a neighbour. If this is 0, they can work without paying for childcare at all. If it is 1, then they pay for childcare according to the age and number of children, and depending on whether the work is full-time or part-time.

Childcare costs depend on whether children are at the age of pre-school, nursery, or school age. If they are of pre-school or nursery age, then part-time work halves the amount

of childcare needed, halving the cost. Meanwhile, in the model those who work part-time do not have to pay for childcare for children of school age.[2]

It is assumed that everyone in the model pays for the cheapest type of childcare, so there is no variation in childcare quality. This is done so as to capture fully the question of how the additional need to pay for childcare affects the labour supply decision of mothers. As the focus of this paper is on the labour market effects of childcare subsidies, rather than developmental outcomes of children, the choice to pay for higher quality childcare, rather than pay for the cheapest form of childcare, can therefore be treated as consumption in this model.[3] As a result, only childcare that is necessary for the a mother to work the number of hours she does is considered in the model.

Treating this additional expenditure as consumption rests on the assumption that parents do not have some threshold for quality only above which they would be willing to use paid childcare - or that this threshold is not above the quality of childcare provided by the median childcare provider in the UK. If this assumption does not hold, we might expect the estimates of these models to slightly underestimate the role of childcare costs on labour supply - as some parents would need to receive more than just median-quality childcare in order to return to work.

1.2.3 Maximisation Problem

The state space for women is $\Omega_{i,t} = \{t, A_{i,t}, \mu^F_{i,t}, h_{i,t}, \mu^M_{i,t}, m_{i,t}, age^K_{i,t}, n_{i,t}, Paid_i, \mathcal{O}_{i,t}\}$. Here, $m_{i,t}$ is a binary variable indicating whether a woman is married, $n_{i,t}$ is the number of children they have, and $age^K_{i,t}$ is the age of their youngest child. $\mu^F_{i,t}$ and $\mu^M_{i,t}$ represent the persistent productivity shocks of the woman (F) and, if applicable, their partner (M). $\mathcal{O}_{i,t}$ represents a binary variable, which is 1 in period t if the individual has faced an unemployment shock in that period, and cannot find work, and 0 otherwise.

The value function of the woman at age t is therefore given by:

$$V_t(\Omega_{i,t}) = max_{q_{i,t}}\{u(P_{i,t}, c_{i,t}) + \beta E_t(V_{t+1}(\Omega_{i,t+1}))\} \tag{1.5}$$

where $q_{i,t} = \{P_{i,t}, c_{i,t}\}$ represents the choices an individual makes each period. The expectation encompasses the probability that a woman will be married next period, the probability that they will have a child, and the possible distribution of their and their partners future wages.

[2] In the case of school age children, childcare costs refer to the cost of before-school and after-school care that would be needed by mothers working full-time.

[3] The choice to use childcare when you have access to non-market childcare, or when the mother does not work, can also be treated as consumption in the model

Within-period utility is given by the following function, for both single and married women:

$$u(P,c) = \frac{(\frac{c}{k}e^{\eta P})^{1-\gamma}}{1-\gamma} \tag{1.6}$$

where k represents an equivalence scale dependent on marriage and the number of children, so to represent the portion of household consumption that benefits the woman.[4] γ can be interpreted as the level of risk aversion, and η can be interpreted as the utility cost of work. η is determined by a non-parametric function:

$$\eta_{i,t} = \eta(m_{i,t}, n_{i,t}, age^K_{i,t}) \times g(P_{i,t}, m_{i,t}) \tag{1.7}$$

where

$$g(P_{i,t}) = \begin{cases} \psi_m & \text{if } P_{i,t} = 0.5 \text{ and } m_{i,t} = 1 \\ \psi_s & \text{if } P_{i,t} = 0.5 \text{ and } m_{i,t} = 0 \\ 1 & \text{otherwise} \end{cases} \tag{1.8}$$

that is, the utility cost of work is dependent on marital status, number of children, age of youngest child, and whether work is part-time or full-time.[5]

1.2.4 Wage Processes

The female wage process is determined by the following equation:

$$\log(w^F_{i,t}) = log(w^F_{initial}) + f^F_i + \log(h_{i,t}) + \mu^F_{i,t} \tag{1.9}$$

so wages are driven by a combination of human capital and productivity shocks, as well as initial heterogeneity f^F_i, which is normally distributed with mean 0. There exists a minimum wage in the model. Thus, if $w_{i,t}$ is determined to be lower than minimum wage, the individual nevertheless earns the minimum wage if they work. The human capital accumulation process

[4] k is set equal to 1 for single women. 0.4 is added for each child, and 0.6 is added if the woman is married. This follows the results of Adda et al. (2017)

[5] It is worth emphasising that the utility function includes the multiplying of utility cost of work by labour supply, so a value of $\psi_m = 1$ would mean that the exponent in the utility function was half as big in the part-time case as in the full-time case.

is given by:

$$\log(h_{i,t}) = \log(h_{i,t-1}) + \alpha_0 I(P_{i,t-1} = 1) + \alpha_1 I(P_{i,t-1} = 0.5) + \alpha_2 I(P_{i,t-1} = 0) - \delta t \tag{1.10}$$

Through this mechanism, the labour participation decision affects the level of human capital, and so impacts future wages. A lower bound is placed on human capital, so that $h_{i,t}$ cannot decay beyond its initial value, which is the value of every individual's human capital when they first enter the workforce. Meanwhile, the persistent shock $\mu^F_{i,t}$ is given by:

$$\mu^F_{i,t} = \nu^F \mu^F_{i,t-1} + \zeta^F_{i,t} \tag{1.11}$$

where $\zeta^F_{i,t}$ is normally distributed with mean zero.

The male income process of the spouse of woman i of age t is given by:

$$\log(w^M_{i,t}) = log(w^M_{initial}) + f^M_i + \gamma_1 t - \gamma_2 t^2 + \mu^M_{i,t} \tag{1.12}$$

where $\mu^M_{i,t}$ is determined analogously to $\mu^F_{i,t}$. For simplicity, the wages of the two parents are uncorrelated.[6]

1.2.5 Exogenous Variables

Women face three other types of shocks, in addition to the persistent productivity shock. The first of these is an unemployment shock. In each period, there is a probability of being unemployed which is correlated only with age, and is not autocorrelated. If women face an unemployment shock, they receive unemployment benefits, and do not face a labour supply decision in that period, as they cannot work.

Women also face shocks to their marital status and fertility, which are both determined exogenously. The probability of being married in at age $t+1$ is given by:

$$Pr(m_{i,t+1}|m_{i,t}, t) \tag{1.13}$$

while the probability of having a child at age $t+1$ is given by:

$$Pr(k_{i,t+1}|m_{i,t}, n_{i,t}, t) \tag{1.14}$$

[6]Were spousal wages positively correlated, women with higher earnings would, relative to where spousal wages were uncorrelated, face higher disincentives from working. This might have knock-on effects for how childcare affects the relative labour supply response of higher- and lower-earning mothers, with knock-on effects for long-run wages.

so the likelihood of getting married or getting divorced is conditional on age, while the probability of having a child is conditional on marital status, number of children and age.[7]

One might have expected women to respond to childcare subsidies by adjusting their fertility - either in timing or on the extensive margin - and thus have to change their labour market response by less, which this model would not capture. However, the literature suggests that there is no significant response of fertility to childcare subsidies (Low et al., 2018; Nollenberger and Rodríguez-Planas, 2015), and so I do not consider this margin.

Finally, the value of $Paid_i$ is exogenously determined for each woman. This is revealed when women have their first child, and is constant from then on.

1.2.6 Solution

The solution of the model consists of policy functions for consumption and labour supply conditional on age, earnings, family demographics and each other element of the state space. There are no analytical expressions for these policy functions. Therefore, I solve the model numerically, beginning with the consumption decisions from age 65 onwards, and then iterating backwards, solving for the consumption and labour supply decision in each period. A detailed description of the solution method can be found in Section A.1 of the Appendix.

1.3 Data

I use data from ten panels of the Understanding Society panel survey, taken from 2009 to 2018. Understanding Society is a representative sample of UK adults with data on a wide range of demographic and labour-market information, which interviews respondents annually. I restrict my sample to women aged between 20 and 65, who did not attend university. Graduate women are excluded for simplicity, as their education decision might interact with fertility choices, adding an additional level of complexity to the model. Graduate women face a very different wage progression to non-graduate women, and thus we might expect them to be more affected by time out of the labour market, and so one might expect the results of this paper to be smaller than the equivalent effects for graduate women.[8]

[7] Marital transitions do not depend on number of children, though there is reason to believe marital transitions, particularly divorce, would be correlated with the presence of children in the household (Tartari, 2015). Accounting for this negative correlation between children and divorce probability might mean mothers would perform less precautionary saving when they had children, and so work less - though it is not clear that this relationship would be affected by the cost of childcare specifically.

[8] Goldin et al. (2024) finds evidence of larger child penalties for graduate women in the long run, possible further evidence of greater labour market scarring for this group.

The survey contains data on their job status, employment history, income, income of their spouse (if applicable), and children and their ages, as well as on their use of childcare. Women who are self-employed, and those who only appear in one panel, are dropped from the survey. Due to their not being accounted for in the model, those with more than three children and those with severe disabilities are also removed from the data. Because they are also not accounted for in the model, and to reduce measurement error, I remove those whose imputed hourly wage is less than the minimum wage. This leaves an unbalanced panel of 17,522 women, with 88,982 yearly observations in total. Statistics on family makeup of those in the data are presented in Table 1.1, and statistics of employment rate of women in each group can be found in Table 1.2.

Number of Children	Single	Married	Total
No Children	0.253	0.105	0.358
One Child	0.118	0.142	0.261
Two Children	0.080	0.199	0.279
Three Children	0.025	0.078	0.103

Table 1.1 Distribution of Family Makeups, Age 20-40

Source: Understanding Society data from 2009 to 2018.

Number of Children	Married	Single	Total
No Children	0.756	0.810	0.792
One Child	0.798	0.666	0.747
Two Children	0.719	0.476	0.655
Three Children	0.471	0.349	0.444
Adult Children	0.703	0.757	0.720

Table 1.2 Women's Employment by Family Makeup, Age 20-65

Source: Understanding Society data from 2009 to 2018. "Adult Children" applies if a family's youngest child is over 15 years old.

Data on employment experience is taken from Understanding Society's employment history module, from which data on total years spent working can be constructed. Data on women's income is calculated in the form of an hourly wage, using data on annual labour income and on hours worked per week. Because there is no labour market participation choice for men in the model, male income is measured on an annual basis, and combines all sources of income.[9] While this paper refers to marriage, the term is also used to apply to those who are cohabiting in the data. Data on childcare usage in Understanding Society

[9]I adjust all income for inflation, adjusting to 2011 price levels.

consists of a question to those who work, what type of childcare they use, and how many hours of that childcare they use.[10]

Understanding Society does not have data on the cost of childcare or expenditure on childcare. Therefore, I obtain data on the cost of childcare directly. For this, I use the Childcare Cost Survey from the Daycare Trust (2011), which was a survey of daycare and out-of-school club costs in local authorities across Britain.

1.4 Estimation

The parameters of this model are estimated in three different ways. Some are set using standard values from the literature and other sources, some are estimated directly from the Understanding Society data, and the rest are estimated using Simulated Method of Moments (SMM).

1.4.1 Externally Set Parameters

Some parameters are not taken from the data or the model, but from external sources. The coefficient of relative risk aversion γ is set to 1.5, following much of the literature, including Blundell et al. (1994) and Attanasio and Weber (1995). I also follow Attanasio et al. (2008) in setting the discount rate β to 0.98 and the interest rate r to 1.5%.

The benefits received are taken from UK legislation in 2011. The benefits accounted for are Income Support, Child Benefit, Working Tax Credit, Child Tax Credit and Jobseekers Allowance, an unemployment benefit. The working tax credit includes a childcare subsidy for children who are 3 or 4, which is also included. [11] Each of these can be calculated with information on the variables included in the state space, including those that constitute income, facts about children, assets and employment status.

1.4.2 Directly Estimated Parameters

Childcare Costs

Childcare costs are estimated from the Childcare Survey 2011. This gives the average weekly cost of childcare for nursery for children under 2, children over 2, and out-of-school clubs. From this, annual costs of childcare are imputed. The cost of childcare is scaled up according

[10] For the purposes of the measurement of moments, I treat someone working full-time with a child under 5 using more than 20 hours of childcare a week as using full-time childcare.

[11] This is the main existing childcare subsidy in the UK. In it, 70% of childcare costs within this age are included in the base elements in the Working Tax Credit, which are tapered away as income increases.

to the number of children, with a cost of childcare 1.5 times higher for women with two children relative to having one child, and 2 times higher for women with three children.[12]

Marriage Probabilities

Marriage probabilities are estimated from the Understanding Society data. In order to generate a sample representative of the 10-year snapshot of the population in the data, the marriage probabilities and divorce probabilities are generated to match the frequency of divorce transitions in the data, and the proportion of married people in 5-year age groups in the data. From this, implied marriage probabilities are derived.

Fertility Probabilities

Each household can have up to three children, and the probability of having a child is taken from frequency of childbirth given certain characteristics in the same data. Factors taken into account are 5-year age group, marital status, and number of children. The resulting simulated dataset closely matches the proportion of single and married women in each age group with each number of children.

Unemployment Probabilities

The probability of becoming unemployed is also taken from the data. Unemployment probability is allowed to vary with age, and is taken from the proportion of those in the labour force who are unemployed in each age group.

Male Wage Process

In the model, it is assumed that men always receive some income, so I do not correct for selection. This income may be from labour income, benefits, or some other source. It is estimated as annual earnings.

The values of γ_1 and γ_2 are estimate by a linear regression of age and age squared on log annual earnings. The residuals are used to estimate the other parameters, which are σ^2_{fM}, ν^M and $\sigma^2_{\zeta M}$, through GMM, taking into account measurement error.

[12] Childcare for an additional childcare costs less than for the first child for two reasons: the first is that it can be cheaper to enter multiple children into the same childcare system. The second is that the state space does not include the ages for older children, and so an additional child may be older by enough that they require cheaper childcare.

The results indicate a high variance of permanent shocks, at 0.039, with fairly low levels of persistence, with ν^M at 0.6. Initial heterogeneity is also high; the initial variance of log annual income is 0.159.[13]

Initial Conditions

Some women have children, have a partner, and have worked before the age of 20. To account for this, I estimate initial conditions based on women aged 20 in the data. The probability of being married is estimated in combination with the probability of having one or two children, as well as the distribution of youngest children. The distribution of number of years worked before age 20 is also taken from the data, and corresponding levels of human capital are estimated for each point in this distribution.

1.4.3 Estimation by Simulated Method of Moments

The remaining parameters are estimated using SMM. I minimise:

$$\min_{\Pi}(\hat{\phi}_{data} - \phi_{sim}(\Pi))' \mathcal{F} (\hat{\phi}_{data} - \phi_{sim}(\Pi)) \tag{1.15}$$

The vector Π contains the rest of the unknown parameters. This includes the disutility of work for different marital status, number of children, and age of youngest children, which make up the possible outputs of the function $\eta(m, n, age^K)$. The adjustment factors made for the utility cost of part-time work, ψ_m and ψ_s, are also included. Π also includes the parameters determining the female wage process. This is both those determining the productivity shock, ν^F and $\sigma^2_{\zeta^M}$, those that determine experience accumulation, α_0, α_1, α_2 and δ, and the others that determine wages, $w^F_{initial}$ and $\sigma^2_{f^F}$. Finally, the proportion of women who do not have access to non-market childcare, λ, is also estimated by SMM.

The empirical moments $\hat{\phi}_{data}$ are calculated using the Understanding Society data. I use moments that relate to wages, extensive and intensive margins of employment, marital status, number and age of children, and childcare use. Simulated moments ϕ_{sim} are calculated by running the full model, and obtaining equivalent statistics from the simulated data. The weighting matrix $\mathcal{F}$ is a diagonal matrix, containing the inverse of the standard deviation of each moment. These standard deviations are calculated from the data, using the bootstrap method.

The moments used can be divided into four categories. There are those that are used to extract the utility cost parameters, those that inform the human capital accumulation process,

[13] Further details of this estimation process, and the estimation of marriage and fertility probabilities, can be found in Section A.2 of the Appendix.

income moments that define other income parameters, and a single moment to draw out the proportion of women who have access to non-market childcare.

Utility Cost Moments

The utility cost moments derive from data on employment rates. Most of the moments derive from a probit regression of a dummy variable representing different values of marital status, number of children, and age of youngest child on employment, controlling for age. The coefficients for each value of the dummy variable are used as moments. Since single childless women are used as the baseline in the probit regression, I also include a moment representing the mean employment rate of single childless women in their late 20s, as a baseline to help determine the levels of utility cost of work across the whole population. Finally, to differentiate the utility cost of part-time and full-time work, I use four moments that represent the ratio of full-time workers to part-time workers, divided by marital status and age group.

	Married			Single		
	Age of Youngest Child			Age of Youngest Child		
	0-5	5-10	10-15	0-5	5-10	10-15
1 Child	-0.76	-0.57	-0.68	-1.56	-1.95	-1.85
	(0.04)	(0.02)	(0.05)	(0.03)	(0.02)	(0.02)
2 Children	-0.91	-0.52	-0.47	-1.62	-1.44	-1.41
	(0.02)	(0.02)	(0.01)	(0.04)	(0.03)	(0.02)
3 Children	-0.98	-0.52	-0.71	-1.41	-1.17	-1.06
	(0.02)	(0.01)	(0.03)	(0.05)	(0.03)	(0.02)
No Children		-0.84			-2.17	
		(0.03)			(0.05)	
Adult Children		-0.70			-1.72	
		(0.01)			(0.01)	
Multiplier for Part-Time Work		0.949			1.369	
		(0.016)			(0.013)	

Table 1.3 Utility Cost of Work Parameters by Number of Children, Age of Children and Marital Status

Standard errors listed under parameter estimates in brackets.

Table 1.3 gives the estimates of parameters relating to the utility cost of work. Note that more negative values mean that work is less attractive. For the part-time multipliers, higher

values mean that part-time work is less attractive relative to full-time work, after taking into account the reduction in hours worked.[14]

The key takeaways from these estimates are that the utility cost of work is higher when single than when married, and that the utility cost of work without children is higher than with children. The intuition for these findings is as follows. Married women have access to some household income, regardless of whether or not they work, whereas single women do not. However, their employment rates are fairly similar, implying a greater non-pecuniary cost to work for single women than for married women.[15] Similarly, the rate of employment among mothers, given the increased pecuniary costs of working with children, is higher than one would expect it to be if their utility cost of work was equivalent to that of non-mothers, implying a higher utility cost of work. Furthermore, in general, the utility cost of work is decreasing in number of children. This implies that childcare costs and government benefits such as the child tax credit and income support can be taken as the drivers of reduced labour market participation from mothers, rather than decreased preferences for work.

Income Moments

To estimate the wage process, I re-specify the wage process outlined in Eq. (1.9) to include mean-zero measurement error. This means the wage process equation is now as follows:

$$\log(w^F_{i,t}) = log(w^F_{initial}) + f^F_i + \log(h_{i,t}) + \mu^F_{i,t} + \epsilon_{i,t} \tag{1.16}$$

This does not change the decisions made by agents in the model, as the measurement error does not affect them, but instead changes only the observations of income in the simulated data. Its variance is then estimated from the real data. This allows for the capturing of other relevant parameters using moments which will be affected by measurement error in the data. Since women can work part-time or full-time in each year in the model, wages are measured in the data at an hourly rate, net of tax, and then annualised.

The income moments are a combination of means, variances and covariances. They include the mean log annual income of 20-year-old women, in order to ascertain the initial wage level. They also include the standard deviation of log annual income at five-year intervals from age 20 to age 60. Finally, they include the autocovariance of log annual income, the standard deviation of log income growth rates, and the autocovariance of log

[14] It should be remembered that the utility cost of work is not the only thing to affect the mental cost of work: because consumption and labour are non-separable, the equivalence scale of consumption will also affect the marginal cost of work. However, this does not change the comparison between types significantly.

[15] The explanation for what this smaller utility cost parameter for married women might be capturing are explored in Jones et al. (2015).

income growth rates. In combination, these identify the parameters determining the wage process, which are the variance and persistence of shocks, the size of measurement error, and the variance of initial heterogeneity in wages.[16]

Parameter		Estimate	Standard Error
Initial log wage	$\log(w^F_{Initial})$	9.425	0.006
s.d. of wage shocks	σ^F_ζ	0.121	0.005
s.d. of initial heterogeneity	σ^F_F	0.164	0.005
Persistence of wage process	ν^F	0.784	0.009

Table 1.4 Income Process Parameters

The estimates are shown in Table 1.4. They show reasonably large fluctuations in income, as well as large initial heterogeneity in income. However, the wage process is fairly persistent.

Experience Moments

The experience moments consist of the mean of log annual earnings at five year intervals of ages over 20. They also include parameters from two regressions. The first of these is the regression of total years worked and total years worked squared on log annual earnings, controlling for fixed effects. The second is the regression of a dummy variable denoting whether work was full-time or part-time on wage growth in the next period, controlling for previous income level. Together, these help capture the relative impact of working on future human capital, and to differentiate the impact of part-time work from full-time work.

Parameter		Estimate	Standard Error
Returns to full-time work	α_0	0.0311	0.0003
Returns to part-time work	α_1	0.0159	0.0004
Returns to not working	α_2	-0.0416	0.0038
Decline in experience from age	δ	0.00084	0.00001

Table 1.5 Human Capital Parameters

The estimates of the human capital accumulation parameters are presented in Table 1.5. They show a significant difference in the returns to full-time work and part-time work, and a much larger gap between the returns to part-time work and not working. This indicates that women have a strong incentive to participate in the labour force each year in some form, so as to avoid significant human capital decay. They also show a significant decline in wages with age.

[16] I estimate the standard deviation of the measurement error to be 0.114.

Childcare Moment

The final moment relates to the proportion of women who do not have access to non-market childcare. Understanding Society only has data on non-market childcare access among working mothers, which can not be considered representative of all mothers, since some might only choose to work because they had access to non-market childcare. Therefore, I use the proportion of mothers in work with access to non-market childcare to identify the proportion of mothers in the population with access to non-market childcare. I do this by estimating the proportion of mothers working full-time with children under 5 who use more than 20 hours of childcare per week.

Parameter		Estimate	Standard Error
Probability of having to pay for childcare	λ	0.483	0.020

Table 1.6 Paid Childcare Probability

The proportion of women without access to non-market childcare is given in Table 1.6. The result implies that just over half of women have access to some other form of childcare. This is despite the fact that amongst mothers of young children working full-time, only 28% use childcare. This indicates that there are very large selection effects from having to pay for childcare, suggesting that those who do not have access to non-market childcare are heavily disinicentivised from working full-time.

1.4.4 Fit of the Model

The results of the model matches the moments used well, and aligns with other statistics not used for estimation that are relevant to the role of childcare costs in the female labour supply decision and the child penalty.

The moments that derive from a probit of categorical variables including marital status, number of children and age of your youngest child on employment match the data very closely, staying within two standard deviations of the data. The human capital regressions are matched similarly well by the model, staying close to the value in the data in each case. The variances and autocovariances of wages and wage growth also fit the data well. Due to the precise nature of the estimates of the first and second order income moments across the life cycle in the data, the confidence intervals are fairly small, so in some cases the data strays outside them, but stays fairly close. Finally, the proportion of working mothers using childcare in the model is matched almost perfectly to the data. A full outline of the moments used and their values in the data and the model can be found in Section A.3.1.

To further understand how well the model fits the data, I follow Honore et al. (2020) and evaluate the significance of each moment in estimating each parameter by evaluating how the standard error of the parameters change if a given moment is removed from the weight matrix. This allows me to find out which moments are driving the precision of the parameters. I find that the employment probit is very important for the utility cost parameters: in most cases, its removal would lead to the variance of the parameter estimate multiplying by a large magnitude. I also find that the ratios of full-time to part-time work are key to the estimation of the part-time work multipliers. I find that the precision of the wage process estimate is dependent on both the variance and autocovariance of wage growth, and on the progression of the variance of wages over the life cycle, while the persistence is mostly driven by the life cycle. Similarly, the returns to experience from work are driven by both the experience regressions and mean wages over the life cycle, while the decline in experience from age is mostly driven by the life cycle. Finally, the childcare use moment pins down the precision of the estimate of the proportion of women with access to childcare. Each of these runs in line with the intuition of the SMM estimation.[17]

The model also fits with key facts about the data that are not represented by the moments in the model. In order to compare the model to the data, I produce a measure of the "child penalty" in earnings that women suffer from not having children. It resembles that used by Kleven et al. (2019), measuring intensive and extensive margins of employment, wages, and total earnings, but I focus on the comparison between mothers and women without children, rather than the comparison between mothers and fathers that they use.

The use of this less common version of the child penalty is necessary here is fathers' wages are exogenous, and so are not a useful benchmark to which to compare maternal wage. Meanwhile, non-mothers are not directly affected by the birth of a child, and so are a more useful benchmark.[18] To the extent that there is selection into childbirth, we might expect women to be more or less likely to have children based on their anticipated earning trajectory, and household income and wealth. This can go in both directions, with those earning more likely to delay fertility until later, while those who never have children much more likely to be those with lower household wealth. We should therefore not expect the benchmark non-mother group to be identical in everything but their propensity for children, but they nevertheless serve a useful comparison.

The consequence of using this measure of the child penalty rather than the father/non-father gap is that the composition of the earnings drop is not split evenly between wages, the

[17] A full table of the relevance of significant moments for certain parameters can be found in Section A.3.2 in the Appendix.

[18] This also allows us to abstract from the so-called fatherhood premium, in which fathers' earn more upon having a child (Bensnes et al., 2023).

intensive margin of employment and the extensive margin. Instead, it is driven by a large drop in employment, both on the intensive and on the extensive margin. Meanwhile, due to selection bias, those who are in employment among mothers have a higher average wage than non-mothers in employment.

I demonstrate this using a subsample of the Understanding Society data, comprised of those who appear in every panel of the sample. I follow their labour market outcomes in the years just before and after having children, and compare them to the outcomes of women of the same age who have not had children. I use this data as a test of the fit of the model on a set of outcomes that are not used in the estimation of the model. [19]

Fig. 1.1 shows the fit of the model to the data, and also demonstrates the short-run change in labour market outcomes that occurs upon becoming a mother. In each case, the model matches the data very closely. Both sets of lines demonstrate that, when measured relative to non-mothers, the child penalty in the short term is driven not by a drop in wages, but by a drop in labour supply. In fact, the model and the data both show an uptick in wages of working mothers relative to non-mothers after childbirth. This is likely driven by selection biases, as those who remain in the workforce upon having children are likely to have higher returns to working, due to having higher wages. Meanwhile, there is a significant drop in both the intensive and extensive margin of labour supply, with hours worked and employment dropping by more than 20%, in both the data and the model. In combination, this adds up to a large drop in real earnings relative to non-mothers in the data, which is captured very well by the model.

The Understanding Society panel only covers ten years, so the child penalty in the data is not shown any further in the above graph. In the Results section below, I use the measure above on the simulated data and extend it further into the future, allowing me to see how the child penalty develops over a longer period, as well as to examine how childcare cost subsidies can change the magnitude and shape of the child penalty.

1.5 Counterfactual Results

I use the estimated model to consider different policy scenarios, and their impact on employment, wages, and welfare. The main comparison I use is to a scenario where the government subsidises childcare costs completely, and funds this subsidy using a proportional tax on labour income. This allows me to see how the reduction in the cost of working for many

[19] The graph representing employment ratio bears some resemblance to some of the moments used, as they include the employment rates of mothers of young children and non-mothers, but there are no moments which represent the wages, hours worked or total earnings of mothers.

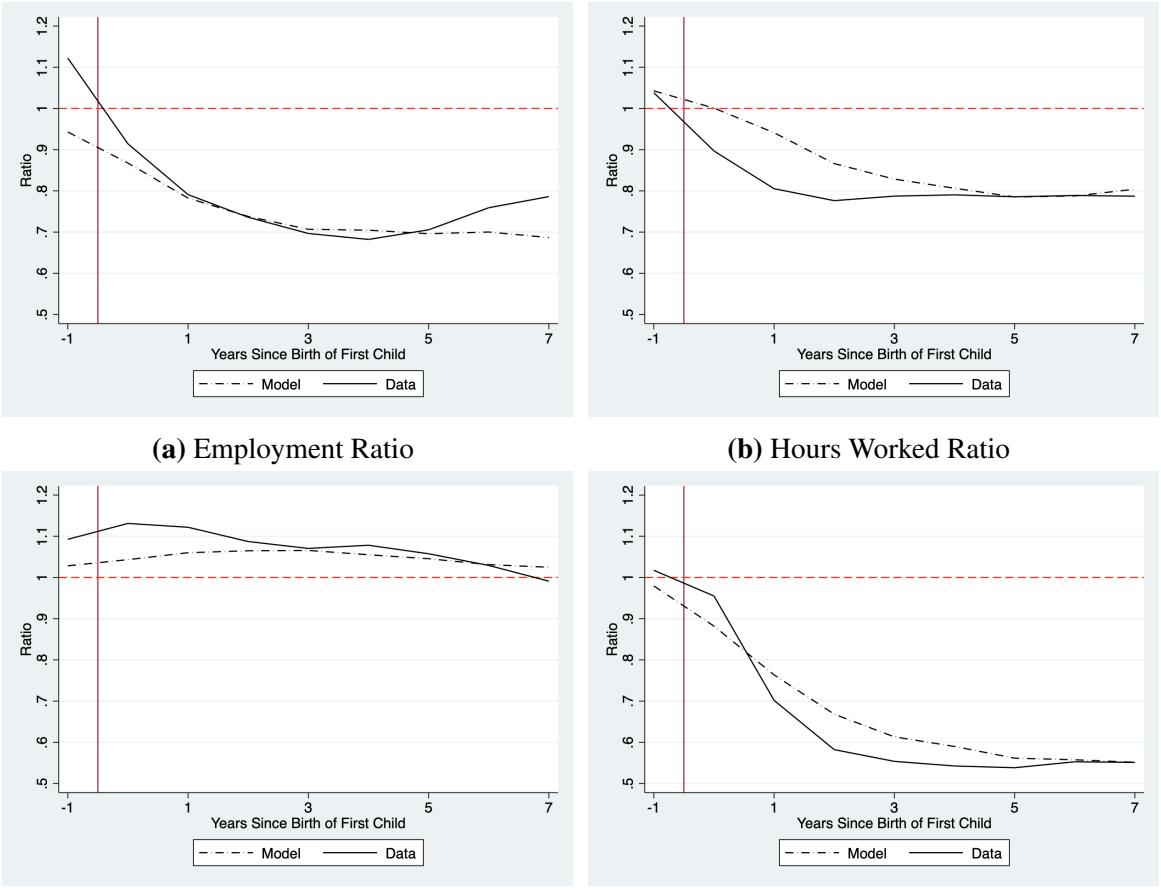

(a) Employment Ratio

(b) Hours Worked Ratio

(c) Wage Ratio

(d) Total Earnings Ratio

Figure 1.1 Ratio of Mothers' Labour Market Statistics relative to Non-Mothers of the same age in the model and the data

Source: Understanding Society. For lines marked Data, simulations for estimated model for lines marked Model. Sample in Data is those who appear in every wave of the Understanding Society dataset and who have a child in the first three waves of the sample. Points are a moving average of three years. The red line indicates the point at which the first child is born. Hours Worked and Wage are conditional on Employment.

young mothers affects their labour supply choices. The fact I model human capital accumulation allows the model to capture the long-run impact of childcare subsidies on wages, as well as on labour supply beyond the point of having young children.

1.5.1 Lifetime Labour Market Outcomes

First, I consider the impact of revenue-neutral childcare subsidies on labour market participation and log wages over the life cycle. The childcare subsidy I consider involves childcare costs being reduced to zero at all stages of early life, including daycare, nursery

and after-school clubs. Therefore, those without access to other forms of childcare do not face any childcare costs. The cost to the government of this policy is dependent on take-up of the subsidy amongst those who *do* have access to another form of childcare, for which they do not pay. The model does not consider preferences for childcare, but it cannot be ruled out that those who previously used a grandparent or neighbour for childcare when nursery was expensive, might choose to use a nursery if it were free. This is because it would not be practically possible for childcare subsidies to be dependent on access to other forms of childcare.[20]

Therefore, I obtain the higher and lower bounds of this problem by considering the size of a proportional tax necessary to fund the childcare subsidy in two cases. The first of these is where no mothers with access to non-market childcare use the childcare subsidy, and the second is where all working mothers use the childcare subsidy, regardless of access to other forms of childcare. In the first case, the proportional tax required to make the policy revenue neutral is 0.3%, in the second case it is 1.4%. This tax is levied on all men and all women in the model, and the reductions in required benefit payments from the introduction of the childcare subsidies are factored into revenue-neutrality. To ensure that the subsidy does not involve any redistribution towards women in the model - so as not to impact the welfare analysis - I do not apply the tax to anyone not in the model, including single men. Therefore, the tax is only applied to non-graduate women and their spouses. In the figures below, I outline both the low-tax case and the high-tax case, to demonstrate the bounds of the policy, depending on the uptake of the subsidy amongst those not already paying for childcare.

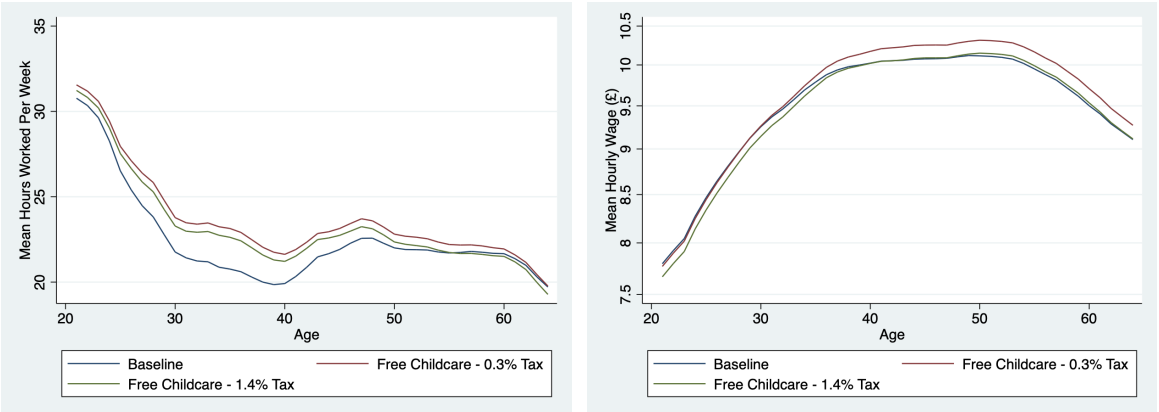

(a) Mean Hours Worked Per Week **(b)** Mean Post-Tax Log Hourly Salary

Figure 1.2 Labour Supply and Income Over Time

Simulations from the estimated model. Data is a moving average of three years.

[20] I implement the condition that childcare subsidies are only available to those who work, either part-time or full-time. Therefore, I do not consider the possibility that women who do not work use subsidised childcare, which would change the cost to the government again.

Fig. 1.2 demonstrates the impact. It shows that childcare subsidies would have a positive impact on both women's labour supply and, in the low-tax case, their wages. Labour supply is shown on the left, as mean hours worked per week, combining the intensive and extensive margins of labour supply. The impact on labour supply occurs mostly early in working life, coinciding with when women have children, by making it cheaper to have children at a young age. The increase in labour supply is over 10% for women in their thirties at either level of uptake. This increase continues beyond the point of having children, but fades out as women reach their fifties.

Wages are shown net of the proportional tax used to fund the childcare subsidy. In this case, the policy has a positive effect on wages only in the low-tax case, mostly occurring later in life. This is driven by the role of increased experience that comes from higher labour supply during the early years of motherhood, and its positive impact on wages over the long run. By the age of 45, this amounts to an average 1% increase in wages in the low-uptake case, which continues on for the rest of working life. In the high-uptake case, the policy leads to approximately a 1% decline in post-tax wages. Some of this difference in post-tax wages is mechanical, and the rest is driven by the reduced incentive effects earlier in life driving down long-run wages.

The change in labour supply and wages is a result of those who do not have anyone who can look after their children in an informal setting. The labour supply of those who do have someone to help with care is unaffected by a childcare subsidy, because although they may use the subsidy if they do work, it does not alleviate any pecuniary cost of working with young children, as they did not necessarily have such a cost anyway. Therefore, I next demonstrate the impact of the same counterfactual subsidy on only those women who do not have access to non-market childcare. The proportion of mothers is one of the parameters estimated by simulated method of moments, and amounts to about 48% of women who have children.

Fig. 1.3 shows that the impact on these women is a lot higher than in the general population. The positive impact on their employment rate is over 25% in their late twenties, depending on the uptake of the subsidy and so on the tax rate. This drives a larger wage effect, reaching between a 4% and an 5% increase in salary from age 45 onwards. This, in turn, keeps a positive employment effect well beyond the point of having children needing childcare.

The results are therefore very different if you consider the proportion of women who need to use market childcare. In the sample of non-university graduate women in the UK, I estimate this proportion to be fairly low, at just under half of mothers. However, Fig. 1.3 demonstrates how large an impact this has on these women, so much so that there is also a

large positive impact on all working-age women, as shown by Fig. 1.2. This indicates that were the proportion of women who did not have a grandparent or friend who could look after their child while they worked were to increase, the impact of childcare subsidies would increase drastically.

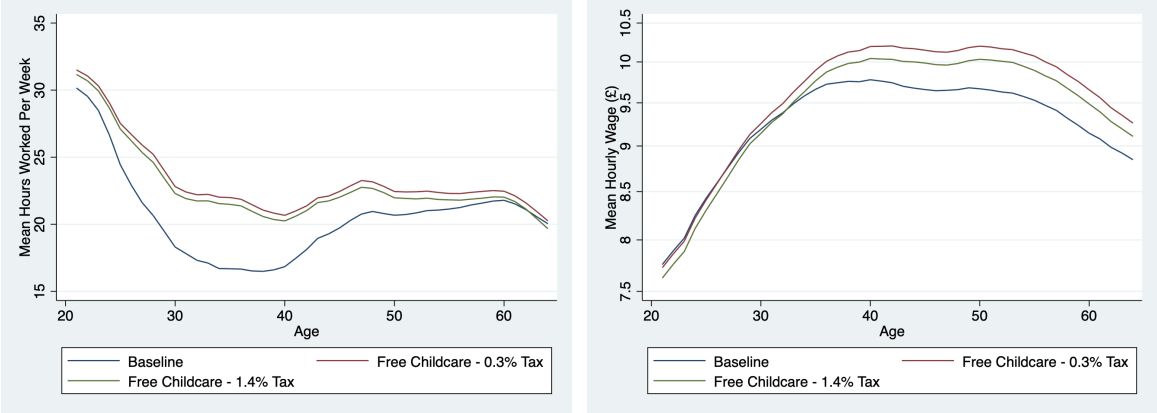

(a) Labour Market Participation **(b)** Post-Tax Log Hourly Salary

Figure 1.3 Employment and Income Over Time Amongst Those Who Have to Pay for Childcare

Simulations from the estimated model. Data is a moving average of three years.

1.5.2 Child Penalty

I also consider the role of the counterfactual childcare subsidy on the child penalty faced by mothers. As in the model fit of Fig. 1.1 I continue to measure this as differentiated from the labour market outcomes of women who do not have children. This is preferred to a comparison to the father of the children, as is sometimes used (Kleven et al., 2019), because the model does not account for male labour supply, so any counterfactual impact on men's labour market outcomes would not be captured by the model. Instead, I consider the former comparison, and extend it to longer-term outcomes than can be measured from the data. This allows me to examine the longer-run labour market penalties of having children, and consider, in the aggregate, how childcare costs impact them.

I consider four labour market outcomes: the extensive margin of employment, the intensive margin of employment, the wage of those working, and labour earnings, the product of the first three. In each case, I present the development of the child penalty over the course of the twenty years after the birth of a woman's first child, and include the counterfactual childcare subsidy outcomes. Because an increase in the tax rate affects both mothers and non-mothers, the impact of different tax rises on the outcome is small, and so I only include the low-tax case. The results are shown in Fig. 1.4.

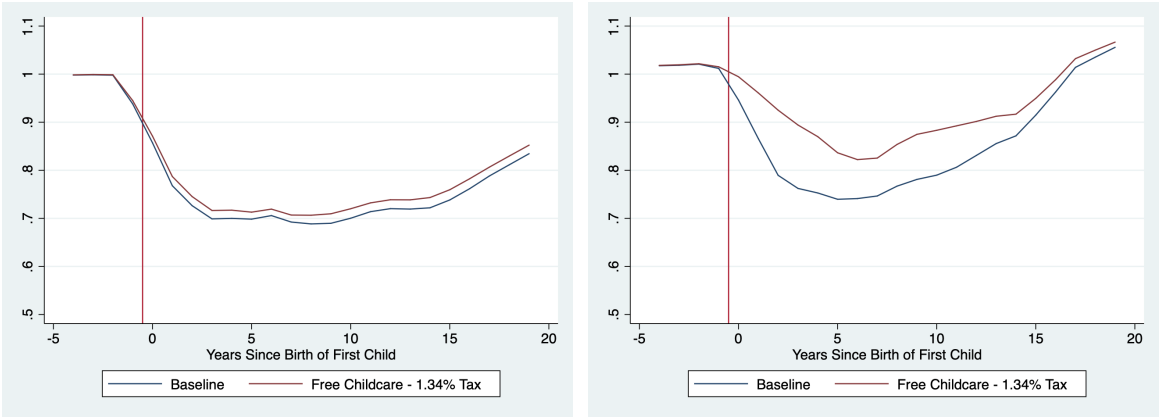

(a) Employment Ratio

(b) Hours Worked Ratio

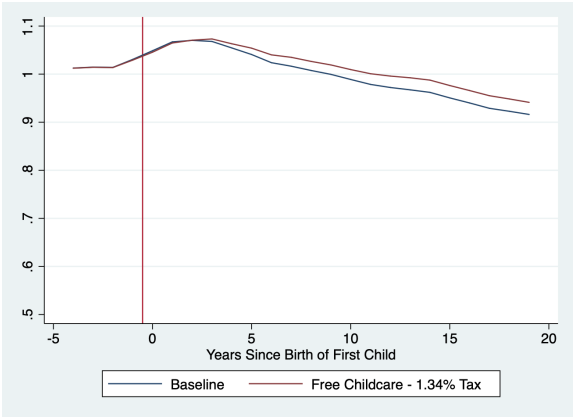

(c) Wage Ratio

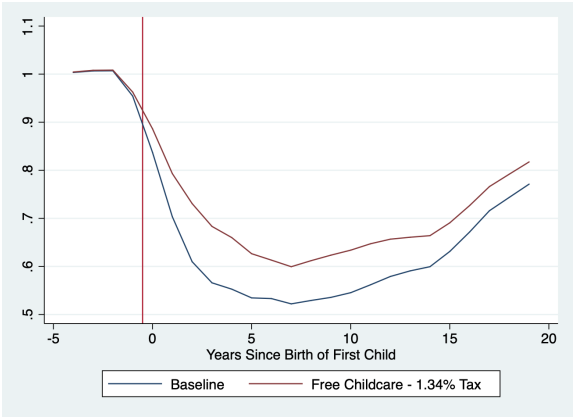

(d) Total Earnings Ratio

Figure 1.4 Ratio of Mothers' Labour Market Statistics relative to Non-Mothers of the same age

Simulations from the estimated model. Points are a moving average of three years. The red line indicates the point at which the first child is born. Hours Worked and Wage are conditional on Employment.

These results give an indication of the size of the child penalty over the long run, and of how a childcare subsidy would change it. The first figure shows the gap in employment between mothers and non-mothers. It is clear that this is the main driver of the total earnings gap, particularly in the first years after childbirth. This gap reaches as high as a 30% drop in employment for much of the first decade after becoming a mother, before gradually shrinking. This graph shows that childcare subsidies would have only a very small impact on reducing this gap. It shrinks the gap by between 1 and 2 percentage points throughout motherhood, representing a reduction between 5% and 10% of the child penalty in the baseline case. The effect continues on over the long run, which can be explained by increased wages among mothers due to accumulated experience.

At the intensive margin of labour supply, the effect of the subsidy is much larger. The results indicate half of the child penalty in hours worked in the first few years of motherhood would be reduced with a childcare subsidy, with the effect deteriorating but not disappearing in the long run. This indicates that the reduction in the cost of working with young children is most strongly felt amongst those working part-time in the baseline case, who are enabled to work full-time by the reduced cost of childcare. This may be explained by the big difference in human capital outcomes between not working and working part-time, which means that those who do not work in the baseline case would need to be incentivised to work throughout motherhood in order for the effect on long-run wages to be worth it. However, the childcare cost subsidy enables those working part-time to shift to full-time work in some years but not others, since their human capital does not decay seriously in either case. Since full-time work has higher returns to experience than part-time work, this effect perpetuates itself, with higher wages driving more hours, driving up future higher wages.

Another explanation for the shift being concentrated in hours lies with the parents of children in primary school. At that age, with childcare costs, there is no cost of childcare for those working part-time, but there is a cost of childcare for those working full-time. So amongst that group, those who were constrained by the cost of childcare would already be working part-time. Since they are the ones constrained, the benefit of a childcare subsidy amongst mothers of older children would display itself in an increase in hours worked, not in entrances to the labour market.[21]

The wage ratio itself follows a very different trajectory. As demonstrated in the data, wages of mothers are higher than wages of non-mothers in the immediate aftermath of childbirth. In the model, this can be explained by selection effects, as the drop in employment means that mostly those with higher wages stay in the labour force, and thus earn higher wages on average than non-mothers. However, this effect peters out within a decade, and the effect of motherhood on wages becomes negative, as more mothers return to the workforce, reducing the selection effect, and the effect of lower human capital amongst mothers kicks in. By twenty years after the birth of their first child, the wages of mothers are on average almost 10% less than those of non-mothers. The childcare subsidy decreases that gap by a third, indicating a large effect of the childcare subsidy on the long run wage penalty of having children.

These three labour market outcomes combine to create the child penalty in terms of total earnings. In the baseline model, this is driven in the short and medium-run by the drop in employment and hours worked. However, in the long run, it can also be attributed to the

[21]This latter explanation cannot fully explain the difference, however, as there is a large change in hours worked among mothers less than five years after the birth of their first child.

wage penalty of having children. The cumulative gap between total earnings of mothers and non-mothers is almost as high as 50%. The childcare subsidy has a significant effect in reducing it: in each of the first 15 years after becoming a mother, the childcare subsidy reduces the total earnings gap by between 20% and 30%, which is a large impact on the child penalty. Beyond this point, when the main effect of the childcare subsidy is the small long-run impact on wages, the impact of the subsidy becomes smaller, but remains a significant reduction in the child penalty relative to the baseline case. This demonstrates that a childcare subsidy would go a long way to reducing the disparity in labour market outcomes between those women who have children and those who do not.

There is some question of what else constitutes this child penalty besides childcare costs. This seems noteworthy given the results of the estimation in Table 1.3 indicate that the utility cost of work decreases upon having children, meaning that the decrease cannot be explained in an increased non-pecuniary cost of work upon becoming a mother. In the model, the explanation instead lies in the benefits system. There are three benefits that drive down employment upon childbirth. The child tax credit and child benefit provide a source of income besides labour upon having children, and are not dependent on employment. Furthermore, Income Support is given only to those who do not work, but is more limited in scope, applying only to single mothers of children under 5. In combination, these benefits provide an income stream that reduces the incentive to enter the labour force, particularly for those on lower incomes. Therefore, even without childcare costs, there is still some child penalty in income over both the short and long run, although significantly decreased relative to current childcare pricing.

1.5.3 Welfare Analysis

The final analysis of the policy counterfactual where childcare costs are eliminated is to consider the ex-ante welfare effects of the policy. I follow Low and Pistaferri (2015) and measure this by considering what percentage change in their annual consumption in the baseline case an individual would be willing to trade for the policy. Because the childcare subsidy is revenue neutral, and because the cost of making it revenue neutral is not exactly estimated in this paper, I provide the welfare analysis for both the low and high uptake bounds of the childcare subsidy. I measure welfare in the first period of the model, so before most people know whether or not what their family circumstances will be, or whether they will have access to non-market childcare. I focus on the median estimate of the welfare effects, but they are fairly uniform across wage heterogeneity and experience.[22]

[22]The only significant heterogeneous effects are for those who have already had children at age 20, for whom the benefits and costs of the policy are very dependent on if they currently need to use childcare.

In the low-uptake, low-tax scenario, where taxes are only raised to cover the subsidy only for those who cannot use any other type of childcare, the childcare subsidy is welfare-decreasing, and has the same ex-ante value as an annual 0.5% increase in consumption for the rest of their working-age life. In the case of higher uptake, requiring a higher tax rise to fund it, the impact of the policy becomes equivalent to an annual 0.9% decrease in consumption for the rest of their working life.[23] This follows logically, as the difference in tax rises in the two scenarios is similar in magnitude to the difference in consumption value of the two scenarios.

This indicates that the direct welfare effects of this policy are to some degree dependent on its uptake, and the cost of funding it that would come with higher uptake. For the policy to be welfare-improving in all cases for non-graduate women, some redistribution towards them would have to occur in the funding for the policy. In all likelihood, any childcare subsidy would not be funded solely by taxes on women and married men, and thus it is possible that a real childcare subsidy might be welfare-improving for women, but not for men.

1.6 Conclusion

This paper considers the broad dynamic implications of a childcare cost subsidy in the UK, both in its effects on short-run labour supply decisions and on long-run labour market outcomes. The life-cycle model demonstrated in this paper allows for the estimation of unobservable variables regarding the utility cost of work, the accumulation of human capital, and access to non-market childcare. By obtaining an estimate of all three of these processes, each of which is a key determinant in the efficacy of a childcare subsidy, I can use the life cycle model to estimate the dynamic effects of a childcare subsidy.

I find that the subsidy does have a significant impact on labour market outcomes, through the mechanism of making it easier for mothers who do not have access to any form of other childcare besides market-rate childcare to work. This creates an 11% increase in female labour supply for women in their thirties, a large increase. This increased labour market participation has some pass through onto long-run wages, increasing post-tax wages by 1% in the long run. There is a stronger effect for the specific group of women whom it affects: a 31% increase in labour supply and a 5% increase in wages. These effects are hindered by the large increase in taxes on non-graduate women used in the model. In a scenario where a childcare subsidy was treated as a redistribution towards women, so that taxes were raised on

[23] It is worth noting that the welfare response in the high-uptake case does not take into account any welfare gains of families switching from non-market childcare to government-subsidised childcare.

everyone and not just non-graduate women, the increases in female employment and wages would likely be higher.

This paper also shows that a childcare subsidy would have an effect on reducing the child penalty, which I measure using a comparison between mothers and non-mothers of the same age. In particular, it would prompt a large reduction in the child penalty on the intensive margin of labour supply, by allowing a lot of mothers who previously worked part-time to work full-time. This, along with a small reduction of the penalty on the extensive margin of employment, has two consequences. The mechanical effect is to significantly reduce the child penalty in total earnings due to an increase in total hours worked. The second is to increase human capital accumulation, and so reduce the smaller child penalty in wages, thus extending the reduction in the child penalty into later life.

2 Adult social care funding and its interaction with saving for retirement

Abstract

This paper examines how the expected cost of social care later in life affects households' decisions on savings and retirement. I use data from the English Longitudinal Study of Ageing on social care use, health, and assets, to estimate individuals' social care risk, and their attitudes towards how to retire. I find that current UK social care funding provides significant insurance against social care risk, and that further expansions of social care funding would not induce significant changes in precautionary saving for retirement, or age of retirement. This is because the current system insures households against the major risk, that of very low consumption in old age, and other saving motives, such as bequests, mean that further precautionary saving is not necessary. I also find that the nature of the social care funding model should have major consequences for the assets households use to save in old age, particularly with regards to the choice between housing and other assets.

2.1 Introduction

Social care in England is a major cost in later life: care homes cost over £30,000 per year, and even small amounts of care at home can cost thousands of pounds per year. The question of how households account for the risk of these high costs when preparing for retirement is therefore an important one, as is the question of how the government insures households

against these risks. I examine the role of elderly social care and its funding, separated from medical expenses, in driving saving and retirement decisions. I look at the consequences for the amount saved by retirement, the makeup of households' asset portfolios, and when they choose to retire. I look at how adjustments to different elements of the frameworks for the funding and provision of social care for the elderly can change how individuals optimise their retirement strategies.

To do this, I build a detailed model of individuals' health and social care risk, and incorporate it into a life-cycle model of households' retirement and saving decisions. This allows us to understand two different questions about elderly social care provision: how highly is it valued, and how does the prospect of high social care costs in old age affect individuals decisions on how to retire? The use of a structural model allows me to estimate the utility individuals receive from social care, enabling me to understand the welfare implications of increased state provision of social care. I incorporate a dynamic element to this model to examine the consequences of social care risk in the long run as well as the short run.

A range of key parameters are necessary to answer these questions. I estimate the processes that govern individuals' need for social care, and incorporate them into the English social care funding system. I use this to estimate individuals preferences for social care use, as well as their preferences for retirement, housing, and bequests. The latter parameters govern how households choose to retire, in terms of when they retire, how much they choose to save for retirement, and the type of asset they choose to save in.

I use this to estimate how large of an effect social care risk has on households' retirement decisions. This is done across several dimensions, considering both the generosity of the social care funding system, and how the social care asset test judges housing wealth, to think about a range of retirement choices. These included the amount households save for and during retirement, what assets they use to save, and when they choose to retire. I use the tests of generosity to estimate how well the current social care funding system insures households against risk, and use different considerations of housing as wealth to estimate the extent to which portfolio allocation is driven by social care funding. I also consider the welfare effects of changing the generosity of the social care funding system, to more explicitly capture how government funding can insure against social care risk.

I use survey data from the English Longitudinal Study of Ageing (ELSA) to estimate the model. I use data on subjective and objective health measures, assets, and social care use to estimate individuals' risk of needing elderly social care in England, and combine the institutional setting of the English social care funding system to estimate individuals' financial risk from social care need. I use ELSA data on assets, housing, and labour supply to

estimate the unobserved parameters that govern how individuals make retirement decisions, using Simulated Method of Moments. Using these parameters, I simulate the same model but with different counterfactual scenarios.

The results show that the current social care funding model plays a major role in insuring households against social care risk. Removing almost all social care support would increase assets at age 70 by approximately 30%, while expanding social care funding to be universal would only lead to a 3% reduction in how much households save, indicating that the current social care funding system prevents the need for most of the precautionary saving brought on by social care. This difference can be seen in two different channels: a difference in savings rate, and a difference in retirement age. Nevertheless, I find that expanding social care support to all individuals would be welfare-increasing, even if the policy were revenue-neutral, indicating that further state support would provide even better insurance.

I also find that housing choices are heavily dependent on the cost of social care. The question of when housing is considered to be an element of wealth that is included in the threshold for receiving help with social care is very important for determining the proportion of household wealth that is held in housing. In particular, I find approximately 5 percentage point increases and decreases in home ownership rates depending on whether the measurement of housing as wealth is more lenient or less lenient than the current policy.

This paper draws on and contributes to several different strands of the literature around saving for retirement and social care. I contribute to the life-cycle literature around saving for retirement, and choices of when to retire. There is a long literature studying saving for old age and the retirement decision (Engen et al., 1999; Scholz et al., 2006). There are various different elements of saving for retirement that the literature engages with, including housing portfolio choice (Crawford and O'Dea, 2020), asset testing for government support (Hubbard et al., 1995; Sefton et al., 2008), and medical need (French and Jones, 2011; De Nardi et al., 2021).

There is also a broader literature on social care, or long-term care, that focuses on its role in saving for retirement (Barczyk and Kredler, 2018; Barczyk et al., 2023; Bueren, 2023), as I do in this paper. Other papers do similar, with a focus on the welfare consequences of insurance for long-term care or healthcare in old age (Kopecky and Koreshkova, 2014; Nardi et al., 2016). Meanwhile, Dobrescu (2015) incorporates the distinction between formal and informal care in household saving. My paper incorporates a range of considerations included in many of these papers, using data from the English social care system. By studying this system, this paper considers the interaction between saving for social care and asset testing. I test the role of policy contingent on care need, wealth, and housing on how households save for retirement.

The rest of the paper is as follows. Section 2.2 outlines the institutional setting of social care in England. Section 2.3 outlines the model, and details how it captures the shocks individuals face, and the decisions they make. Section 2.4 outlines the data used, while Section 2.5 outlines the estimation process, of both the parameters captured directly from ELSA data and by Simulated Method of Moments. Section 2.6 demonstrates the results of the counterfactual analysis, while Section 2.7 concludes.

2.2 Institutional Setting

In the England, a distinction is drawn between healthcare and social care. Healthcare constitutes medical care, and is provided free at the point of use for all individuals, funded by taxation. Social care consists of assistance with everyday tasks, which can range from food shopping to help washing and eating, depending on an individual's need. Unlike healthcare, social care is not free at the point of use. Instead, whether social care can be received free of charge is determined by an individual's assets. In England, those with savings of over £23,250 are expected to pay for their own care. This asset figure does not include the value of someone's home, unless there would be no-one living there if they were to go into care. Effectively, this means that housing is only included in the asset test if every member of the household needs full-time residential care.

If your assets are below £23,250, your care is organised by the local council. They undertakes a 'needs assessment' that adjudicates not just your ability to pay, but also the amount and type of social care you need. If the council pays for the social care, then this may be by the council assigning you a personal budget to be spent on social care, or by the council organising care for you themselves.

2.3 Life-Cycle Model

I model the saving, retirement, and social care use decisions of households from age 50 until death. These households can comprise of couples, single men, or single women. Households make decisions on how much to spend, both on consumption and on social care, as well as whether or not to retire if they are still working, in response to their previous decisions, and to health and social care shocks they experience. The saving decision also accounts for the possibility of saving money in housing, or saving in non-housing assets. This structure allows me to consider how different methods of social care funding would affect households retirement planning across multiple axes: when to retire, how much to save, and whether to save it in housing or in other assets.

Households are modelled from age 50 until death. The age of 50 is chosen due to data restrictions, as it is the starting age of the English Longitudinal Study of Aging (ELSA).[1] Each period t in the model represents households' consumption and retirement decisions at age t. Various cutoffs are put in place in the model for simplicity: households must have retired by age 75, and death is a certainty at age 100. Death occurs stochastically, with its probability determined by age and health. Households with two members are assumed to consist of a man and a woman, with the man two years older than the women.[2] The possibility of divorce is not accounted for.

2.3.1 Problem of the Individual

At age t, households make decisions on how much to spend on general goods, $c_{j,t}$, how much to spend on social care, $CARE_{j,t}$, whether or not to retire, and what house to live in, $Move \in \{0, \frac{2}{3}, 1\}$. $Move$ is a tertiary variable, where the three options are to stay in the same home, downsize to a home worth two thirds of the value of their current home, or to sell their home and begin to rent instead. The model only captures the possibility of downsizing or selling your home completely, rather than any increase in housing, because of the model's focus on social care; in the context of social care, the key housing questions are how it affects your ability to keep your home.

Similarly, the decision to retire is an optimal stopping problem, so that after retirement, it is not possible to work again. In order to capture the joint retirement decision made by many households, married couples make their retirement decision together, and stop working at the same time. [3]

Additional state variables non-housing assets ($A_{j,t}$) and housing wealth ($h_{j,t}$) are a function of the decisions made in the model. Non-housing wealth includes all forms of positive and negative wealth, including cash, private pensions, and mortgage debt. Private pensions are included directly in assets, despite the requirement that they not be spent before the age of 65, because this amounts to a small section of the model. In the case of couples, it is assumed all assets, both housing and non-housing, are jointly held by the couple, rather than by an individual member of the household. $R_{j,t}$ is treated as a state variable, representing whether the household has retired. The first element of the budget constraint for a household

[1] In treating events before 50 as exogenous, we ignore the effect of social care costs on precautionary saving at younger ages, a possible limitation of this work.

[2] When age is discussed in the paper, in the case of married women it is two years older than their actual age.

[3] Where households do not retire jointly in the data, the retirement of the first member of the couple is treated as a decline in wages.

j, with number of members n, at age t is therefore as follows:

$$\frac{A_{j,t+1}}{1+r} = A_{j,t} - c_{j,t} - CARE_{j,t} + w_{j,n,t}(1 - R_{j,t,}) + B_{n,t} + (1 - Move_{j,t})h_{j,t} \quad (2.1)$$

That is, households spend money on general consumption and social care, and receive income in the form of non-housing assets from working, from receiving the state pension $B_{n,t}$, which is a function of age and number of people in the household, and proceeds from any house sale they have made. r represents the risk free return on non-housing assets.

The second element of the budget constraint is the borrowing constraint. The model does not allow for any borrowing besides mortgage debt. This gives the following borrowing constraint for household j at age t:

$$A_{j,t} > -lev(t)h_{j,t} \quad (2.2)$$

$lev(t)$ is the leverage ratio: the proportion of the value of their house individuals are able to mortgage as a function of their age. Since it is only possible to borrow through mortgage debt, the lowest possible level of non-housing assets (including mortgage debt) is the amount they are able to mortgage. In reality, since private pensions are included in non-housing wealth, it is rare for homeowners to come close to this budget constraint.

The final element of the budget constraint is the progression of housing wealth. This is modelled as having a different risk-free return to non-housing wealth, r_h. The progression of housing wealth is given by the following equation:

$$h_{j,t+1} = (1 + r_h)Move_{j,t}h_{j,t} \quad (2.3)$$

2.3.2 Health Process and Death

The health variable of individual i $H_{i,t}$ is a categorical variable, with four levels: 'very good health', 'good health', 'bad health', and 'very bad health'. It plays a key role in determining need for social care, and on the probability of individuals dying. It is governed by a Markov process, where:

$$P(H_{i,t+1} = x) = f(H_{i,t}, M) \quad (2.4)$$

This indicated that future health outcomes are a function of current health and of an individual's gender.

It is possible for individuals to die in this model. If an individual dies, and there is still an individual left in the household, then the remaining individual retains all of the household

assets. If there are no other individuals left in the household, their assets are left as a bequest. I do not account for any knock-on effects of the death of a partner on health, spending behaviour, or labour supply. This is done for simplicity, and because insurance against death is not the key outcome studied in the paper, but it means the model likely misses ways in which individuals self-insure against the death of a partner (Fadlon and Nielsen, 2021). This may mean individuals in the model engage in more precautionary saving for this reason than they otherwise might.

An individual's probability of death is given by the following process:

$$P(Death) = f(H_{i,t}, t) \tag{2.5}$$

This means that the probability of death is a function of an individual's health, and of their age - specifically, the ten-year age bracket they fall into.

2.3.3 Social Care

I estimate the probability that an individual needs social care, and their level of need for social care. 'Need' for social care is determined as the level that is determined by local government when assigning social care to those who are unable to pay for themselves. In the model, this is treated as the minimum level of social care that an individual would ever choose to receive. This is measured in pounds per year, and in the case of local government organisation, is estimated as the number of hours of care someone is assigned, multiplied by the average wage of a care worker.

The probability of needing some social care and the level of need for social care is driven not just by health, but also by age, gender, and past need for social care. The probability of needing some social care ($CARE_{i,t}^{NEED} > 0$) is estimated from the following logit model:

$$y_{i,t}^{1} = \alpha_1 + \beta_1 H_{i,t} + \beta_2 t + \beta_3 t^2 + \beta_4 t^3 + \beta_5 M_i + \beta_6 \mathbb{1}[CARE_{i,t-1}^{NEED} > 0] \tag{2.6}$$

This means that for an individual in the model, there is risk in the model as to whether they will need a positive amount of social care in the forthcoming year, though it is possible to know the exact probability with which they will need a positive amount of social care. This uncertainty conditional on health, age, and past care need, is best interpreted as the uncertain progression of what could be thought of as the "social care need" process, which is caused by and correlated with, but not identical to, health. For example, two individuals with the same health problems might have different need for help with everyday tasks, that can not be predicted from future health status, which is itself uncertain.

Given that an individual needs social care, there are two different equations that determine how much social care they need. The first of these is the possibility that they might need to be in full-time residential care. Within this, no distinction is made between different levels of care, as the total number of hours of care they need is constant at 168 hours per week. The probability that an individual needs full-time care ($CARE_{i,t}^{NEED} > 30000$) is estimated conditional on them having positive care need. This is given by the following logit model:

$$y_{i,t}^{2} = \alpha_2 + \gamma_1 H_{i,t} + \gamma_2 t + \gamma_3 M_i + \gamma_4 \mathbb{1}[CARE_{i,t-1}^{NEED} > 30000] \tag{2.7}$$

This adds an additional layer of uncertainty for individuals in the model. Conditional on needing care, there is uncertainty as to whether that care will need to be full-time, or can be received in an individual's own home.

If the individual does not need full-time residential care, then log weekly hours of care need (z) is estimated using the following OLS regression:

$$log(z_{i,t}) = \alpha_3 + \delta_1 H_{i,t} + \delta_2 t + \delta_3 t^2 + \delta_3 t^3 + \delta_4 M_i \tag{2.8}$$

Care need in pounds per hour is then equal to z multiplied by 52.5, multiplied by the average care worker's hourly salary. This means there is not uncertainty in the level of need for non-residential care, conditional on finding out that some is needed. This reflects that in the data, low-wealth individuals with similar observable characteristics who are receiving care tend to receive very similar amounts of care.

In the model, social care need acts as a lower bound to social care usage. It is possible for individuals to spend more on social care per year than local government would assign to them given their needs. This is true both if they would receive some free care from the government, and if they would not. The model does not discriminate between spending more on more hours of care, or spending more on higher quality care. In the case of those resident in care homes, the margin of further expenditure will be quality, as the number of hours spent in care is fixed. For in-home care, however, it is possible that the main margin of further expenditure will be hours.

It is also possible that individuals spend no money on social care, because they have others who will help them with it. This can be their spouse, or it can be other friends, neighbours, or family members, such as children or siblings, who live nearby. There is therefore a state variable fam_j that represents whether there is someone who is willing to help your household with care. It is assumed that those who have a friend or family member who provide them with care use this for all their care, as it is very uncommon for individuals who receive at least some care from people they know to pay for any care. Furthermore, this

variable is constant over time, so those with someone able to care for them continue to have someone throughout old age, while those who do not, continue not to have someone to care for them. Similarly, it is assumed that if you are married, and your spouse does not need any care for themselves, they provide care for you.[4] Note that this only applies to at-home care, and not to those who need to move into a residential care home - where this is the case, presence of family members makes no difference to social care costs.

Asset and housing test

The model includes the means tests described in Section 3.2, such that individuals only have to pay for their own care *needs* if their wealth falls below certain thresholds. However, this only covers care which they need - they would still be expected to pay for care they might want beyond the assessed level of need.

The main asset test is that household non-housing wealth must be below £23,250 in order for individuals to receive free provision of the level of care they are assessed as needing - if their non-housing wealth is above that, none of their care is provided. Housing is excluded from the asset test - creating an incentive for households to store their wealth in housing - unless the care needed is such that the house would be empty if due to the care (for example, if both members of a couple needed residential care). If this is the case, the value of the house is included in the asset test.

2.3.4 Maximisation Problem

The state space for a household j at time t is:

$$\Omega_{j,t} = \{t, A_{j,t}, h_{j,t}, R_{j,t}, H^m_{j,t}, H^f_{j,t}, CARE^{NEED,m}_{j,t}, CARE^{NEED,f}_{j,t}, fam_j\} \tag{2.9}$$

Here, $H^m_{j,t}$ and $H^f_{j,t}$ represent the health of the man and the woman in the household respectively. If one of them is deceased, or if the individual was always single, then the value of $H_{i,t}$ for that individual is 'dead'. Similarly, in this case, the value of their need for care is 0.

The value function of a household j at age t is therefore given by:

$$V_t(\Omega_{j,t}) = max_{q_{j,t}}\{u_1(c_{j,t}, CARE_{j,t}, R_{j,t}, h_{j,t}) + \tag{2.10}$$

$$\beta((1-\Psi)E_t(V_{t+1}(\Omega_{j,t+1}))\} + \Psi u_2(A_{j,t}, h_{j,t}))\} \tag{2.11}$$

[4] Over 80% of those who receive care from their partner or a family member do not receive any other care. Less than 5% of individuals who receive some care from their partner or a family member receive more than half of their care through paid care.

where $q_{j,t} = \{c_{j,t}, CARE_{j,t}, Retire_{j,t}, Move_{j,t}\}$ represents the choices a household makes each period. There are two possible outcomes in the next period, each discounted by discount rate β: with probability Ψ, all individuals in the household will be dead, and utility will derive from bequests. With probability $1 - \Psi$, there will still be at least one individual left in the household. Their expectation is across all possible health, social care and death shocks that the household could encounter.

Within period utility is dependent on consumption, care expenditure, retirement, and housing. The relationship to utility is dependent on age (in the case of retirement) and on number of people in the household (in the case of housing). The utility function is as follows:

$$u_1(c, CARE, R, h) = \frac{\frac{c}{k(N)o(h)}^{1-\sigma}}{1-\sigma} - \eta_t(1-R) + \alpha_N \mathbb{1}[h > 0] + f(CARE) \tag{2.12}$$

Utility in this model is almost completely separable, with utility from consumption independent of labour supply and care spend. Utility from consumption depends on the equivalence scale $k(N)$, which is set as 1 if there is one individual in the household, and as 1.6 if the household consists of a couple. Utility from consumption also depends on home ownership, with $o(h)$ varying depending on whether the individual owns their house or is renting, to reflect the fact that renters will have to spend some of their consumption on rental costs, which owners would not have to. Utility from consumption has constant relative risk aversion, with risk aversion coefficient σ. Utility from retirement is given by the $\eta_t(1-R)$ term, which reflects that the utility cost of work changes linearly over time, adjusting for the model starting at age 50:

$$\eta_t = \eta_1 + \eta_2(t - 50) \tag{2.13}$$

Thus the utility cost of work at the start of the model is η_1, and it changes by η_2 every year after that.

It is worth noting that the utility function does not depend on health, either through consumption or through the utility cost of work. Instead, health only contributes to risks that determine costs faced by an individual. This is a simplifying assumption due to the difficulty in separately identifying the role of health on consumption and labour supply choices. One might expect precautionary saving to also be affected by a link between health and the utility function - for example, one might need to save less to insure against pure health, because worse health might also lower one's marginal utility of consumption. This is something not considered in the model.

There is also a non-separable term representing utility from housing, $\alpha_N \mathbb{1}[h > 0]$. This represents only utility from owning a home, not any utility that might be obtained from owning a more expensive home. It may be that individuals choose to move to a smaller home when they are older out of preferences, and so we do not apply any change in utility to moving from a larger home to a smaller home. It is worth noting also that the utility from home ownership is dependent on whether a household is a couple or not. Home ownership rates amongst the elderly vary dramatically by marital status: one would expect utility from ownership for couples to be positive, but in the case of single people, particularly widowers, it is possible they may have a preference for selling their home.

The final term of the utility function represents utility from care. This term is given by the following function:

$$f(CARE_{j,t}) = \begin{cases} \epsilon(max_N)^{1-\sigma}(1-\sigma)^{-1} & \text{if } CT_{j,t} = 0 \\ \epsilon(CARE_{j,t}\frac{max_N}{CT_{j,t}})^{1-\sigma}(1-\sigma)^{-1} & \text{if } CT_{j,t} > 0 \end{cases} \tag{2.14}$$

Here, $CT_{j,t} = CARE_{j,t}^{NEED,m} + CARE_{j,t}^{NEED,f}$, represents the total amount of care needed by the household. max_N represents the maximum amount of care that might be needed by a household with number of members N. For a married couple, this represents the cost of two people in residential care. If there is no need for care, households receive a fixed amount of utility from care, and so do not choose to obtain any care. Utility from care is separable from utility from consumption - possible interactions between the two, which might, for example, lead individuals to spend down their savings slower when they were in residential care, are ignored.

If there is need for care, individuals get utility from additional care. The term $\frac{max_N}{CT_{j,t}}$ indicates that households get more utility from an additional pound of care if they have less serious care needs, as less is needed to be spent in order to meet their care needs. It also means that individuals with different levels of care needs all have exactly the same utility from care when they are receiving that amount of care. Utility from care has the same coefficient of risk aversion as from consumption, σ, but is adjusted by ϵ based on households' relative preferences for spending on care relative to spending on consumption of other goods. An additional assumption underlying these expressions is that households with no need for care obtain no utility from care.[5] Similarly, it is assumed that those who only need at-home care would not choose to pay for residential care instead.

[5]This is also important for the identification of care need, as I assume that the fact of being observed to use a positive amount of care cannot be caused by preferences for care.

Utility from bequests is received upon death of the last member of household j between age t and age $t+1$, and is given by the following equation:

$$u_2(A_{j,t}, h_{j,t}) = \zeta \frac{(A_{j,t} + h_{j,t} + 1)^{1-\sigma}}{1-\sigma} \tag{2.15}$$

Effectively, utility from bequests is equivalent to spending all remaining assets, both non-housing and housing, adjusted by some parameter ζ.[6] The same risk aversion coefficient σ applies. No distinction is made between leaving housing and non-housing assets as a bequest.

2.3.5 Wages

This model uses a very simplified model of the labour market. Households containing a married couple have a combined labour income, while single individuals have their own labour income. Labour income progresses deterministically over time, meaning that there are no wage shocks in the economy, though initial wages are heterogeneous. The progression of household incomes over time is given by the following equation:

$$w_{j,t} = w_j^{Initial} + w_1(t-50) + w_2(t-50)^2 \tag{2.16}$$

Thus each household knows its expected wages up until retirement, and there are no income shocks to provide uncertainty. This is done for computational simplicity, as the purpose of this model is examine saving and retirement responses to health and social care risk, rather than to model income risk. However, to the extent that there is income risk for these workers, it may be that they may engage in precautionary saving to protect themselves against these risks, which they might then be able to use for social care spending. Due to this assumption, the model would not capture this.

2.3.6 Solution

The solution of the model consists of policy functions for consumption, care use, retirement, and housing choice conditional on the state space. There are no analytical expressions for these policy functions. Therefore, I solve the model numerically, beginning with consumption, care and housing decisions made at age 100, and then iterating backwards, solving for decisions in each period. A detailed description of the solution method can be found in the Appendix.

[6] Note the $+1$ term in Eq. (2.15) exists only to avoid infinitely negative utility for those who die with no assets at all.

2.4 Data

I use data from nine waves of the English Longitudinal Study of Ageing (ELSA), spanning from 2002 to 2019. ELSA is a nationally representative sample of UK adults over 50 with data on a wide range of demographic, saving, and labour-market information. Respondents are interviewed every two years. I use the whole sample, removing only those with a partner who does not live with them. All nine waves of the survey contain data on marital status, health conditions, self-reported health, job status, labour income, house value, non-housing assets, mortgage debt. From wave 6 to wave 9, there is also detailed data on care receipt. This includes data on number of hours of care per week, amount paid for care, and who provided the care, broken down by individual. Therefore, data on care is taken only from waves 6 to 9, covering 2012 to 2019, while data on labour market outcomes, saving choices, and health is used from throughout ELSA.

Private pensions data is not collected from wave 5 onwards. Data from waves 1 to 4 is therefore used to impute private pension holdings of individuals from wave 5 onwards, using information on retirement, income, savings, and housing, as well as past private pension holdings where possible.

Data on care is much better for those receiving at-home care than for those who move into a residential home. Where ELSA has data on those receiving residential care, it often does not have data on assets, and so data from the previous wave is used. Those in residential care are also under-sampled, and so the data is re-weighted using estimates of the care home population of the UK. Finally, unlike with at-home care, ELSA does not have data on the amount spent on residential care, so this data is not used to inform how much people spend on care in the model. Data on deaths is taken from the ELSA End of Life Survey, which took place in waves 2, 4, and 6 of the study.

2.4.1 Initial Conditions

The model is estimated on a population taken from ELSA. To ensure a large enough sample, data is taken from those households where the representative adult is between 50 and 54. Those with obvious anomalies are removed, and some single women are removed at random to ensure an even gender balance in the initial group - ELSA has a disproportionately large number of women in the sample at age 50 (beyond the imbalance naturally arising from higher mortality in young men than young women). This gives a sample of 1995 households in their early fifties. In each case, we have data on the number of members of the household, the health of each individual, the retirement status of the household, their labour income,

Variable	Mean	Standard Deviation
Log Household Income	10.14	1.766
Log House Value	12.41	0.632
Log Total Assets	13.34	1.16
% Retired	2.3%	
% Homeowners	86.1%	
% Married Households	73.8%	

whether they own a house and its value, their mortgage, and their total non-housing assets. Summary statistics for this group are shown below:

2.5 Estimation

The parameters of this model are estimated in three different ways. Some are set using standard values from the literature or from legislation, some are directly estimated from ELSA data, and the rest are estimated using Simulated Method of Moments (SMM).

2.5.1 Externally Set Parameters

Some parameters are taken directly from the literature. The coefficient of relative risk aversion σ is set to 1.5, following much of the literature, including Blundell et al. (1994) and Attanasio and Weber (1995). I follow Attanasio et al. (2008) in setting the discount rate β to 0.98.

The thresholds and rules for social care receipt are taken from legislation. This includes the asset threshold above which individuals have to fund their own care, set at £23,250, as well as the rules regarding home ownership, and when it counts towards the threshold. The state pension level and eligibility are also taken from legislation - the state pension in the model is set to £8528 per year per person, and the cutoff age is set to 65, as it was in 2019. Care home costs and details on the care home population, used for weighting the data, are taken from the Office for National Statistics (ONS). Meanwhile, data on average pay for care workers is taken from The King's Fund.

2.5.2 Directly Estimated Parameters

Health and Health Process

In order to estimate the health process, it is first necessary to create a good measure of health. ELSA contains two different health metrics: self-reporting of health, and indicator variables

for a wide range of health conditions. Using either of these independently can face pitfalls: self-reporting of health is subjective and can be subject to justification bias (Kreider and Pepper, 2007), while it is difficult to transform individual health measures into a single health index. Since Stern (1989), this problem has been circumvented by using objective health measures as an instrument for self-reported health. I follow the specific method used by Blundell et al. (2017), who carry this out on ELSA data. In the first stage of this process, I run a regression of different health conditions on self-reported health:

$$SubjHealth_{i,t} = a + bX_{i,t} \tag{2.17}$$

Where $SubjHealth_{i,t}$ is an individual's self-reported health, a base 5 categorical variable ranging from 'poor'' to 'excellent', and $X_{i,t}$ is a vector of indicator variables, representing different health conditions. Full results of this regression, and the list of health variables considered, are presented in the Appendix. I use this regression to predict each individual's self-reported health from their health, and use this prediction as the basis of a health index. I then discretise it into 4 categories, matching the self-reported health variable, and use this as the health variable $H_{i,t}$ used in the model.

In order to estimate the probabilities of the Markov process underlying health, I use the longitudinal element of ELSA. I estimate the proportion of individuals of a given level of health, for each possible previous level of health in the last wave, for each gender. This faces the issue that the waves are two years apart, while the model I apply it to is annual. To resolve this, I take the square root of the probability of remaining the same level of health, and use this as the probability of staying the same level of health in one year. I then adjust the other probabilities proportionately. Thus, in order for an individual to remain the same level of health over two years, they probability of not changing would have to be multiplied by itself, to give the figure from the data. This rules out the possibility of individuals experiencing two shocks, the latter reversing the former, in two years, but because health is persistent, this is unlikely to be common, and so this method provides a good estimate of the underlying probability of health transitions.

Social Care Process

As discussed in Section 2.3.3, there are three equations which govern social care use and need in the model. These are the equation governing the probability that an individual will need some social care, the equation governing the probability that an individual will need residential care, and the equation governing the amount of at-home care an individual needs if they do not need residential care.

In order to estimate the equation governing social care need, I run a logistic regression (with $F(.)$ representing the logit function) with the same specification as Eq. (2.6):

$$care^*_{i,t} = \alpha_1 + \beta_1 H_{i,t} + \beta_2 t + \beta_3 t^2 + \beta_4 t^3 + \beta_5 M_i + \beta_6 \mathbb{1}[CARE^{NEED}_{i,t-1} > 0] + u_{i,t}, \quad (2.18)$$

$$\mathbb{1}[CARE^{NEED}_{i,t} > 0] = \begin{cases} 1 & \text{if } care^*_{i,t} > 0, \\ 0 & \text{otherwise.} \end{cases}$$

The probability of needing social care - observed in the data as having a level of care greater than zero - is therefore estimated from an individual's health, age (allowing for curvature by including higher order polynomials), gender and whether they needed social care in the previous period. This regression exploits the assumptions of Section 2.3.4 that no one who needs no care would choose to pay for some, though those who need some might choose to pay for more. If this were incorrectly assumed, the estimates from Eq. (2.18) might imply a higher probability of care need for individuals than was truly the case. I use the point estimates of these regressions to inform the logit regression in the model; this informs both the probability that an individual will need care, and individuals' expectations that they will need care. The results of this regression are as follows:

	Probability of Any Care	
Independent Variable	Regression Estimate	Standard Error
Health	0.706	0.025
Age	0.053	0.003
Male	-0.268	0.051
Past Care Need	3.453	0.064
Constant	-8.211	0.212
N	67,490	

Table 2.1 Probability of Needing Any Care

As these results show, the largest determinant of social care need is previous social care need; it is unlikely that those who currently need social care will not need it in future. Following this, shifting health by one category is more important for affecting social care need than gender, although gender is significant, with men less likely to need care than women. Due to the use of higher order polynomials for age, the effect of age is difficult to interpret, although probability of social care need does increase monotonically with age.

Analogously to the estimation of social care need, I run a logistic regression similar to Eq. (2.19) in order to estimate the probability that an individual needs residential care, once again allowing for measurement error. Higher order polynomials in age are not included

here, as they are not significant:

$$care^*_{i,t} = \alpha_2 + \gamma_1 H_{i,t} + \gamma_2 t + \gamma_3 M_i + \gamma_4 \mathbb{1}[CARE^{NEED}_{i,t-1} > 30000] + u_{i,t}, \tag{2.19}$$

$$\mathbb{1}[CARE^{NEED}_{i,t} > 30000] = \begin{cases} 1 & \text{if } care^*_{i,t} > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Once again, the point estimates of these regressions are used to inform the logit regression in the model, generating the probability that an individual will need care in a care home, and their expectations of needing that care.[7] This regression is run only on those individuals who use some care, and so is used to generate the probability of needing residential care conditional on needing some care. The results of this regression are as follows:

	Probability of Residential Care	
Independent Variable	Regression Estimate	Standard Error
Health	0.325	0.094
Age	0.081	0.012
Past Residential Care	3.389	0.360
Constant	-8.824	1.037
N	5221	

Table 2.2 Probability of Needing Residential Care

As with care need in general, residential care need is strongly positively related to previous residential care need, indicating that those in residential care are likely to stay there. Health plays a smaller role in determining residential care need than general care need, while it appears that age plays a larger role: ageing by five years has a larger effect on the probability of needing residential care than shifting from very good to very bad health.

Estimating social care need for those individuals who do not need residential care is more difficult, as some may choose to spend more of their savings on care. I attempt to capture only those who cannot afford to pay for other types of care. I do this by estimating an analogous OLS regression to the process in Eq. (2.8). However, it is estimated only on those with non-housing assets below £5000 - I use this as a threshold below which it is unlikely individuals will spend their own money on care. Thus the hours of care for these individuals should not reflect preferences, but instead care. The sample is also restricted only to those who need some non-residential care. This gives the following regression:

$$log(CARE^{HOURS}_{i,t}) = \alpha_3 + \delta_1 H_{i,t} + \delta_2 t + \delta_3 t^2 + \delta_3 t^3 + \delta_4 M_i + u_{i,t} \tag{2.20}$$

[7] Analogously to 2.18, this regression assumes that being in residential care is a reflection solely of need, and not of preferences.

This estimates how much care individuals who cannot afford to pay for their own care receive, giving an estimate of care need, as this paper defines it. The results are as follows:

	$log(CARE_{i,t}^{HOURS})$	
Independent Variable	Regression Estimate	Standard Error
Health	0.260	0.126
Age	0.029	0.008
Man	0.508	0.254
Constant	-1.173	0.845
N	151	

Table 2.3 Number of Hours of Care Needed

Health plays a significant role in determining care need, while age plays less of a role - variation in care need, amongst those who do need care, is fairly minor across age.

Informal Care Access

The proportion of individuals who have access to help with care from family, friends and neighbours is taken from ELSA data. Different proportions are used for married households and unmarried households. The assumption that where family members are able to help, they always do, is used. This is justified by the fact that it is very rare for individuals to receive both informal care and professional care.[8] Therefore, the proportion of households with access to informal care is taken to be the proportion of households who need care, and receive it informally, rather than by paying for it - in the case of married couples, this is only relevant when both individuals need care. To estimate the probability that someone will have access to care, I use the proportion of single individual's who receive help with care, as estimates for couples are harder, due to them helping each other. This tells us that quite a large proportion have access to help with at-home care: approximately 71%.

Mortality Rates

The probability of a death in the model is estimated in a similar manner to health transitions. For a subset of individuals who have died, ELSA has data on their health and age before dying. I use ONS data on the mortality rates of different age groups, and re-weight the ELSA data to match these mortality rates. Thus, the effect of age on probability of death comes from ONS data, rather than from ELSA. For each ten-year age bracket, I then estimate the relative probability of dying for each different level of health. This gives death probabilities conditional on age and health.

[8] See the Appendix for more information on informal care use.

Other Parameters

Other parameters are also estimated directly from ELSA and other sources. One of these is the real interest rate r, which is taken from the mean savings rate return over the period considered, subtracting mean inflation. This gives a real return on savings of 1.14%. Another parameter estimated directly, this time from ELSA, is r_h, which is estimated using the mean change, as a percentage, in housing value for those who do not move house. Since the model does not include housing risk, I follow Kaplan and Violante (2014) by subtracting the variance of returns. Due to the difference in length between the periods in the model and the waves in the data, the square root of this value is taken, to get the average percentage change in housing value in one year. This gives a real return on housing of 0.3% over the time period considered from 2002 to 2019. This low real return is likely a reflection of the housing crash in the UK during the Great Recession.

2.5.3 Estimation by Simulated Method of Moments

The remaining parameters are estimated using SMM. I minimise:

$$\min_{\Pi}(\hat{\phi}_{data} - \phi_{sim}(\Pi))'\mathcal{F}(\hat{\phi}_{data} - \phi_{sim}(\Pi)) \tag{2.21}$$

The vector Π contains the rest of the unknown parameters. This includes the parameters determining the base level and change over time of the utility cost of work, η_1 and η_2. It also includes the parameters governing preferences for housing for single individuals and for married couples, α_1 and α_2, as well as the parameters governing the progression of the income process, w_1 and w_2. Finally, it includes the two parameters that govern relative preferences towards bequests and care, ζ and ϵ respectively.

The empirical moments $\hat{\phi}_{data}$ are calculated using data from ELSA. Moments relating to home ownership rates, asset holdings, wages, retirement rates, and the relationship between care use and asset holdings are used. Simulated moments ϕ_{sim} are calculated by running the full model, and obtaining equivalent statistics from the simulated data. The weighting matrix $\mathcal{F}$ is a diagonal matrix, containing the inverse of the standard error of each moment.

Moments can be linked fairly clearly with the category of moments they belong with. These can be divided into income moments, retirement moments, housing moments, a bequest moment, and a social care moment.

Income Moments

Income follows a deterministic path in this model, with heterogeneity taken from the initial population in the ELSA data. Therefore, the only parameters that need to be estimated are those governing the progression of income over time. It is necessary to estimate it using simulated method of moments due to the selection effect of households' retirement decisions, which have an effect on which individuals' wages can be observed.

Income for married couples is considered jointly in this model. To generate equivalence between single individuals and couples, I double the income of single individuals, in both the data and the model, for use in the income moment conditions.

I use the following moments, in each case representing only the results for those who are not retired:

Moment	Value	Standard Error
Mean Log Annual Income, Age 57	10.23	0.02
Mean Log Annual Income, Age 62	9.98	0.02
Mean Log Annual Income, Age 67	9.45	0.04

Table 2.4 Income Moments

These moments represent household income, and so can be thought of as representing the total amount earned by a married couple in a year. The years 57, 62 and 67 are chosen for five-year intervals, and so as to avoid possible sharp changes around the state pension age of 65. These identify the moments governing the progression of household wages over time.

Parameter		Estimate	Standard Error
Linear Log Wage Change With Age	w_1	-0.0020	0.0015
Quadratic Log Wage Change With Age	w_2	-0.0044	0.0002

Table 2.5 Income Parameters

This shows that incomes decrease over time, and that the rate of decrease of wages increases over time. At age 50, households lose about 1.6% of their income each year, a value that gets higher as time goes on. This follows from the moments, which show increasing drops in levels of income, despite the fact that those retiring are likely to have lower incomes.

These results may reflect not just productivity declining with age, but also that those who retire choose to work fewer hours, or more enjoyable but less well paid jobs. Since the model does not account for change along those margins, this is captured in decreasing wages for older workers.

2.5.4 Retirement Moments

In order to estimate the utility cost of working in this model, I use similar moments to those used for income. As in the model, I define retirement in the data as when the entire household is retired. In the ELSA data, that means that every member of the household report themselves as 'Retired', 'Permanently Sick or Disabled', or 'Looking After Home Or Family', though in the last case, only if there is not evidence of them working after that point. This generates the following moments:

Moment	Value	Standard Error
Proportion Retired, Age 57	0.142	0.008
Proportion Retired, Age 62	0.315	0.009
Proportion Retired, Age 67	0.680	0.009

Table 2.6 Retirement Moments

These moments identify the two parameters, which govern the original utility cost of work at age 50, and how that utility cost of work changes over time:

Parameter		Estimate	Standard Error
Utility Cost of Work, Age 50	η_1	0.00580	0.00022
Change in Cost of Work Each Year	η_2	-0.00019	0.00002

Table 2.7 Retirement Parameters

The utility cost of work parameter is difficult to interpret. In terms of consumption, 0.00330 represents approximately the change in utility of a single person increasing their non-housing consumption from £10,000 to £14,000. It is also noteworthy that the utility cost of work decreases as time goes on, perhaps counterintuitively. This can likely be explained by the fact that those who continue to work later in life, particularly above state pension age, are more likely to be doing jobs they prefer. These parameters therefore capture this shift, though it may just apply to those who are on the margins of working or not working at a later age.

A key element here is that utility cost of work in the model is not heterogeneous. This does mean that it misses the increasing utility cost of work for many individuals, but those individuals are ones who choose to retire fairly early anyway. For the purposes of considering social care counterfactuals, the utility cost of work of those making the marginal decision whether or not to work is most relevant, so for that purpose, this utility cost of work measure is reasonable.

2.5.5 Housing Moments

In order to estimate preferences for housing, home ownership moments are used. Because preferences for housing are based on the binary indicator of home ownership, rather than on utility from the value of a home, rates of home ownership are preferred to mean house values. The moments are as follows:

Moment	Value	Standard Error
Home Ownership Rate, Single, Age 67	0.667	0.016
Home Ownership Rate, Single, Age 77	0.637	0.017
Home Ownership Rate, Single, Age 87	0.652	0.023
Home Ownership Rate, Married, Age 67	0.899	0.007
Home Ownership Rate, Married, Age 77	0.879	0.008
Home Ownership Rate, Married, Age 87	0.895	0.018

Table 2.8 Housing Moments

This shows that there is a large gap in home ownership rates between married couples and single people, lasting deep into old age. This gap exists at age 50, and is reflected in the initial conditions of the model. However, it persists even though there are likely to be frequent transitions from married households to single households, as a result of the death of one member of the household. Meanwhile, home ownership rates amongst married couples remain consistently high. These moments generate the following values for the parameters that govern preferences for housing, for both single individuals and married couples:

Parameter		Estimate	Standard Error
Utility from Home Ownership, Single	η_1	-0.0020	0.00005
Utility from Home Ownership, Married	η_2	0.0002	0.00003

Table 2.9 Housing Parameters

These results show that the utility from home ownership is positive for married couples, and negative for single individuals. As is the case for the retirement moments, this likely reflects the fact that utility from home ownership is homogenous across all individuals. It is likely to in part capture a particular negative preference for home ownership amongst the recently widowed, causing home ownership rates among single individuals to drop, despite a large majority of individuals who become single by becoming widowed being homeowners.

2.5.6 Bequest and Care Moments

The moments used to estimate the final two parameters, those that estimate the relative strength of preferences for bequests and care relative to preferences for consumption. In

order to estimate the strength of individuals' bequest motives, I use the average level of total assets at age 85, as an indicator of how much individuals might plan to leave when they die [9]. Meanwhile, the method for generating a moment to identify households' preferences for care is to estimate the average amount spent on care by those who receive some care. This is only done for those who receive at-care, as ELSA does not include data on nursing home expenditures.

Moment	Value	Standard Error
Mean Log Total Assets, Age 85	5.052	0.055
Mean Log Care Spend	8.116	0.059

Table 2.10 Other Moments

These moments give the following parameter values:

Parameter		Estimate	Standard Error
Preference for Bequests	ζ	1.56	0.14
Preference for Care	ϵ	0.41	0.05

Table 2.11 Other Parameters

This shows that bequests are valued over three times more than a year of consumption, indicating that there is a strong bequest motive for individuals. Meanwhile, preferences for care are weaker than those for consumption. It is worth noting that this applies on the margin, given that in the model, individuals always choose to spend on what they would need. Thus their preferences for care *beyond this point* are weaker than their preferences for consumption.

2.6 Counterfactuals

I use these parameter estimates, estimated both directly from the data and using SMM, to construct different counterfactuals. The dynamic element of the model can be used to examine how saving and retirement decisions in the run-up to needing care change with changes to how social care is funded. In particular, by considering households' decisions in the run-up to needing social care, I examine how well the current funding structure insures households against the cost of social care.

To do this, I first consider two different counterfactuals, which are the extremes of possible government funding for social care. In the most generous case, the government

[9] Note that the level of total assets is taken in thousands of pounds, before the logarithm is taken

funds social care costs according to their needs assessment of the household, no matter the size or makeup of the wealth of a household. I define the least generous case to still include funding for social care for those who have no other option at all, so households still receive social care if they have no money at all; effectively, if a household has any assets at all at the end of a period, it is always taken to fund any social care needs that a household may have.

Doing this allows us to see the extent to which precautionary saving for social care risk currently drives saving, and the extent to which current social care funding reduces the need for precautionary saving relative to a much more limited form of government support.

I then consider another mechanism by which social care funding might impact saving for retirement, which is saving portfolio choice. I do this by considering the counterfactual where housing is not treated differently to other forms of assets in the test for social care. This allows me to examine the extent to which social care risk can drive not just the amount households save for retirement, but also through what medium households save.

2.6.1 Comparison of Social Care Funding by Generosity

In order to tell how well households are insured against social care risk by the current social care funding structure, I examine two extremes of generosity. The current system involves receiving help with social care if your assets are below £23,250, in which case the government provides you with a baseline level of social care. The £23,250 figure only includes housing wealth if your care would mean you moving out of your home, and it would mean no one else was living there.

For the maximum generosity case, I assume that there is no longer a threshold which households need to fall below, and so everyone receives their baseline needed level of social care. It is still possible for households to spend money on higher quality or more social care if they choose to, but their baseline need is covered.

The minimum generosity case is a little more complicated. Because the model uses discrete time periods of a year, rather than continuous time, the counterfactual is generated that at the end of the year, those with social care needs and any money must use that money to pay for their social care. Effectively, this means that those with social care needs greater than their remaining assets at the end of the year will end the year with no assets, incentivising households to spend all of their savings and so receive the same level of social care as if they had not.

It is difficult to ascertain exactly how generous the current social care policy is, in terms of the proportion of the total amount households spend on care which is spent by the government. Data from ELSA indicates that for those receiving care in their own home, only 7% of the hours of care received from professional carers is paid for by the government. An ONS report

suggests that approximately 50% of those in care homes due to old age are self-funding (Office for National Statistics, n.d.). This model produces a lower estimate, suggesting that based on household asset predictions - which are matched in the model - approximately 25% of the money spent on care homes should be by the government. The most likely explanation for this gap is households avoiding care home fees by giving large gifts while they are alive, something not accounted for in the model. Nevertheless, the implication is that this model should be thought of as a lower bound for how generous the current policy is, and that accounting for these gifts would increase the quality of the insurance.

In order to examine how well the policies at the two extremes insure households relative to current policies, I compare how much they save in each case. I do this both by comparing the savings rates in each case, and comparing the consequent level of assets as a result of these savings. I also compare the average retirement age in each of the three cases.

Firstly, the comparison with current policy of the savings rate in each case is shown in Fig. 2.1.

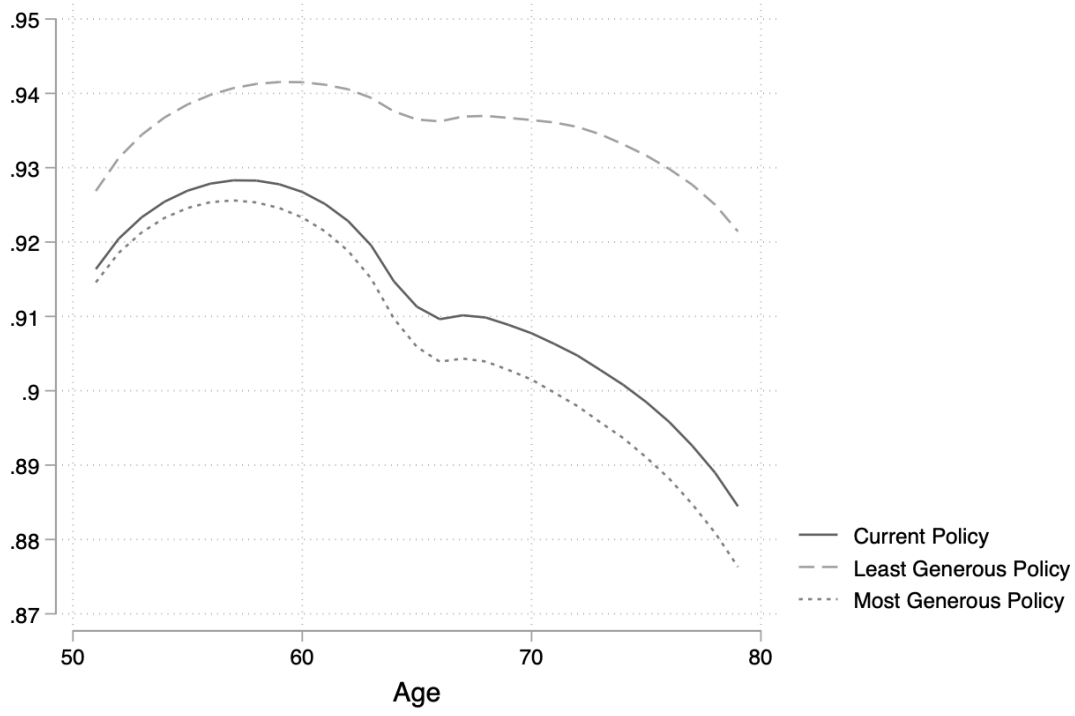

Figure 2.1 Savings Rates Over Time

Simulations from the estimated model. Values are a moving average over a three year period.

This shows that the current funding structure for social case provokes saving responses much closer to the most generous possible policy than it does to the strictest possible policy. By the age of 65, households spend 2.5 percentage points less of their wealth each year than

they would if there was almost no social care funding structure. On the other hand, they only spend about 0.5 p.p. more of their wealth in the current case than they would if everyone had a baseline level of social care funding guaranteed for them.

Similar results are visible in the other margin of saving for retirement, which is when individuals choose to retire. These results are visible in Table 2.12. They show that the average age at which the whole household will have retired is much higher when social care funding is less generous. On the other hand, households do not much work longer under the current policy than they would if they were guaranteed to receive some help with social care.

	Current Policy	Least Generous	Most Generous
Mean Retirement Age	63.8	66.6	63.6

Table 2.12 Average Retirement Age, by Generosity of Policy

Unsurprisingly, the combination of large gaps in savings rates and retirement choices between the current policy and much less generous policies leads to very large differences in levels of total assets in each case. Fig. 2.2 demonstrates this, highlighting a huge gap in assets between the current policy and the low-generosity case, and a very small gap in assets between the current policy and the high-generosity case. Compared with the low-generosity case, assets currently are approximately 30% lower by age 70, and the gap increases as time goes on.[10] This continues deep into old age, indicating that it would be rare for households to actually reach a state of having zero assets, because of the high cost to consumption of doing so. In contrast, there is only a very small drop in assets in the high-generosity case relative to the current policy, about 3%, which is erased by the increase in assets from savings on social care in old age.

This indicates that the main purpose of saving for social care may be to avoid the worst-case scenario, that of completely spending down savings, and then having to continue to pay for social care costs even at that point. Not only does this result in very low consumption in old age, but also removes any bequest, which I estimate as an important motive for saving. On the other hand, it appears that household saving is much less affected by the expected cost of social care above that which will get paid for currently. This may be explained by the complementarity of saving for social care and for other expenditure: households are already saving for retirement for consumption in old age and bequests anyway, and those costs can also contribute towards the cost of care in the case that it is needed.

Therefore, it appears that the current social care model does a very good job of insuring households against social care risk: despite covering at most 50% of the total amount spent

[10] It is worth noting that the model treats the decisions individuals make before age 50 as given, so this is an underestimate of the effect, as it is treated as an exogenous change to social care funding that happens when individuals are 50.

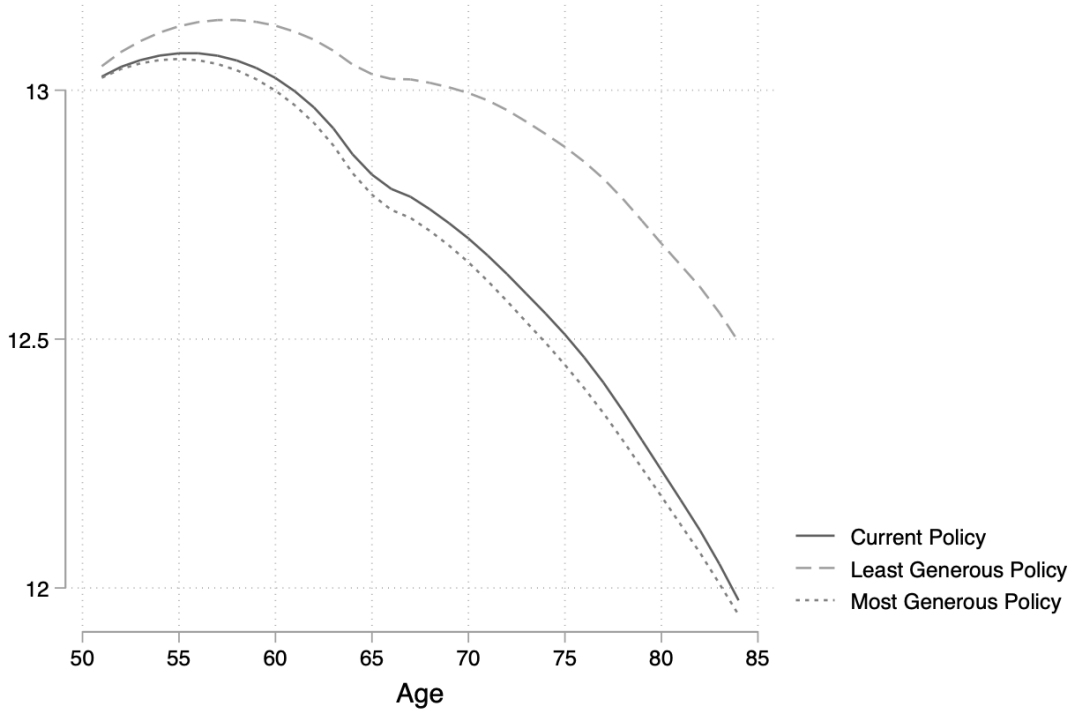

Figure 2.2 Log Assets Over Time

Simulations from the estimated model. Values are a moving average over a three year period.

on social care, it prevents almost entirely the need for precautionary saving against social care risk.

2.6.2 Comparison by Nature of Housing Test

As a second set of counterfactuals, I estimate the role that social care funding plays in the nature of a household's saving portfolio. I estimate this by considering a different counterfactual, where housing is excluded from the asset test. This means that even in the case of everyone in a household needing residential care, housing wealth is not counted towards the asset test. This counterfactual is designed to give a sense of the role of the social care funding model on home ownership.

I measure the effect of a different housing policy with two different measures. The first is the proportion of home-owning households in the two cases, and the second is the proportion of a households total portfolio made up of housing. The first of these, the effect of the nature of this asset test on home-ownership rates, is shown in Fig. 2.3. This demonstrates that the nature of the social care asset test plays a major role in households' portfolio decisions, even before there is a significant risk of losing one's home. By age 70, when very few people have

amassed large social care cost, home ownership is 5 percentage points lower than it would be if housing was never included in the asset test. This indicates that much of the decline in home ownership as households get older would not happen if their house continued to protect them against social care costs.

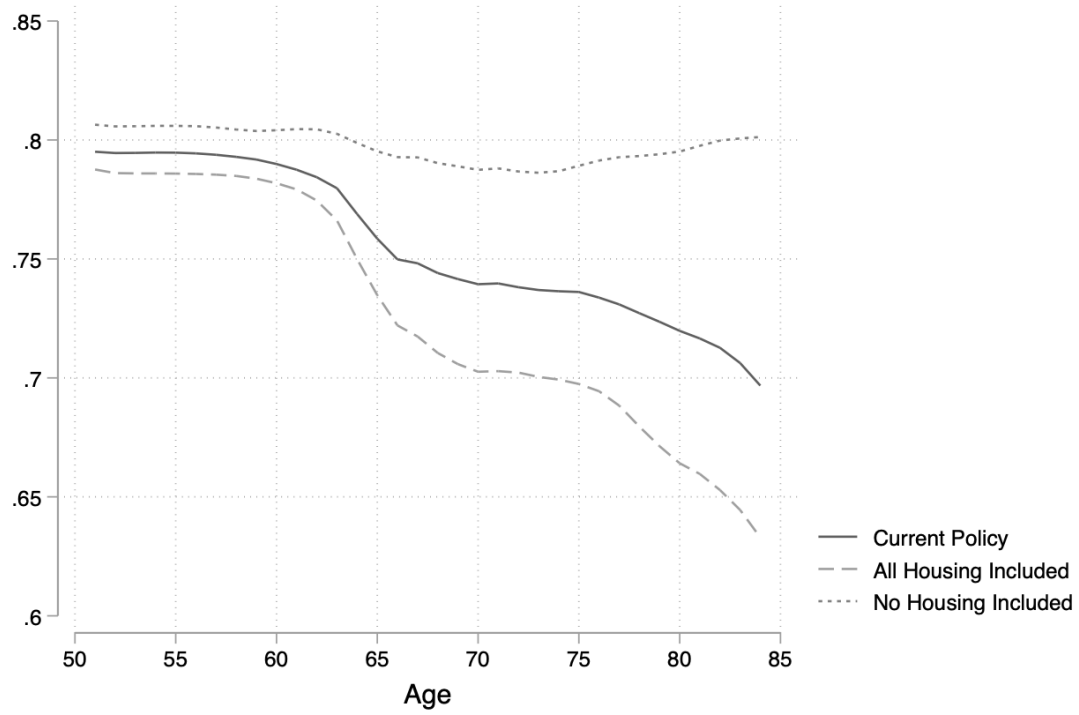

Figure 2.3 Home Ownership Over Time

Simulations from the estimated model. Values are a moving average over a three year period.

Similar results emerge from an analysis of how the social care housing test affects the makeup of a household's asset portfolio. Fig. 2.4 demonstrates how this process varies over time. If no housing were included in the asset test, housing would make up, on average, 10 percentage points more of households' portfolios than it would under current government policy. This indicates that not only are housing decisions changing on the extensive margin, but also that if housing were fully protected from asset tests, fewer individuals would downsize their homes. This makes intuitive sense, as the receipts from downsizing would count towards a social care means test, but if the house is unsold it would not count. What is noteworthy, however, is the size of the effect, and the fact that it happens decades before it is likely that people will have to pay for social care.

The counterfactuals used in this paper demonstrate that the model of social care funding used has significant consequences for how individuals save for retirement. Both in the generosity of the asset test, and in how different assets are treated, the social care means test

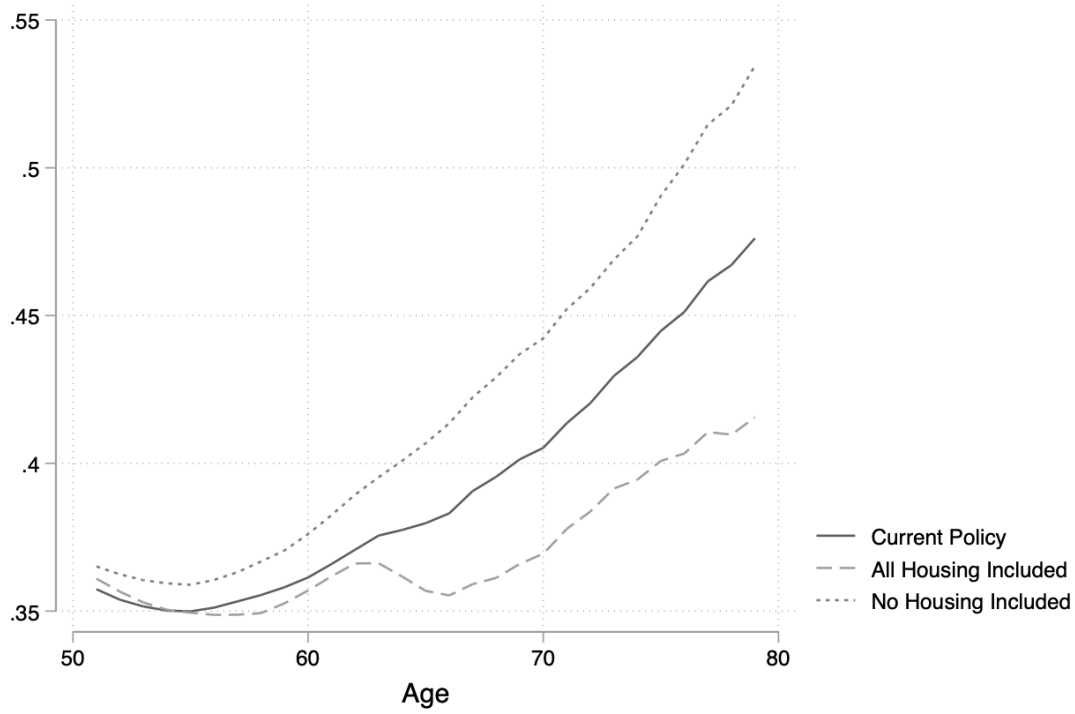

Figure 2.4 Portfolio Allocation Over Time

Simulations from the estimated model. Values are a moving average over a three year period.

affects household decisions leading up to retirement along three different dimensions: how much to save for retirement, in what assets to save it in, and when to retire.

2.6.3 Welfare Analysis

The final analysis of this paper takes a different approach to considering the insurance value of government funded social care. This approach involves estimating the ex-ante welfare effects of the different policy counterfactuals discussed above. Instead of analysing the effect of different policies on saving, I instead follow Low and Pistaferri (2015) and estimate how much additional annual consumption a household would trade for one of the different policies, making the decision at age 50, before they have encountered elderly social care. This more directly captures the insurance values of different social care policies.

In order to avoid capturing the income effects that come from different social care policies, I make each policy revenue-neutral for the government, so it does not involve a net transfer to or from the sample as a whole. I do this by introducing an annual wealth tax on all individuals τ, such that each individual has to pay the government a proportion of their non-housing wealth $\tau A_{i,t}$. For each policy counterfactual, a tax (or subsidy) of this format is introduced

such that the difference between current government social care spending and counterfactual spending is equal to the amount of tax raised. Thus, the population can, on average, only be made better off by the insurance value of a policy, not by positive income effects.

I examine the same policies as before: removing the asset test, making the asset test only apply to those with no assets at all, always including housing costs in the means test, and never including housing in the means test. Table 2.13 lists the taxes and subsidies on wealth required to fund each policy. With respect to the level of the means test, it shows that eliminating the means test would cost as much as severely restricting it would save. Meanwhile, keeping housing permanently in the means test would only need a very small subsidy to be compensated for, for two reasons: firstly, it would only raise money for at-home care costs, which are much smaller than residential care costs. Secondly, as seen in Fig. 2.3, we can expect many people to sell their homes, which would then be followed by spending the revenue from the house sale until they fall below the means test again.

Policy	Tax (Subsidy)
No Asset Test	0.4%
Least Generous Asset Test	(0.2%)
Housing Never in Means Test	0.14%

Table 2.13 Tax Required to Make Each Policy Revenue Neutral

All welfare effects are detailed in Table 2.14. I find that the welfare effects of changing the asset test level are large. Households would value a revenue neutral policy change to remove the means test for social care funding as highly as 6% increase in their annual consumption for the rest of their life. The negative effects of tightening the asset test are even larger. Households would get a similar loss in ex-ante utility from making the asset test as strict as possible as from an 18% decrease in their annual consumption. This comparison demonstrates two things: firstly, as the savings result in Fig. 2.1 also indicates, the insurance value of the current funding setup is far greater than would be provided by any further expansions. However, it also shows that a further increase in the means test would bring large further increases in welfare.

Policy	Equivalent Consumption Change
No Asset Test	6%
Least Generous Asset Test	-18%
Housing Never in Means Test	6%

Table 2.14 Tax Required to Make Each Policy Revenue Neutral

With respect to housing, the more generous approach, of excluding housing from the means test, has similarly positive effects to removing the asset test. While it is a less generous

policy, and so the immediate welfare gains of the policy are smaller, it is also cheaper for the government, and thus requires a lower asset test, which makes up the rest of the difference. Given the relatively small utility effects from home ownership in the model, it seems unlikely that the utility benefit comes from people no longer having to sell their home, but rather from exempting more people from the costs of social care, and providing similar insurance as to the removal of the asset test.

The results of the welfare analysis demonstrate that there would be a major insurance value from a more expansive social care funding system, while also reinforcing the evidence that the current funding system provides a large amount of insurance. It is worth bearing in mind that this analysis assumes that individuals have a good understanding of the social care risks they face. However, ELSA asks respondents about their likelihood of needing certain types of care. Their responses indicate an overestimation of the probability that they will need care:

This raises interesting behavioural questions, given that it is possible to purchase social care insurance in the England, but it is not widespread. This could be explained by a lack of information about social care risks, or about the degree of current government coverage. It could also be a function of greater discounting of specific risks.

2.7 Conclusion

In this paper, I have considered the dynamic implications of social care costs, both in their broad influence on saving and in the specifics of the English social care funding context. By estimating the process of need for social care, and by using Simulated Method of Moments to estimate household preferences for different saving drivers and for social care itself, I am able to estimate the interaction between social care risks and saving choices for those of retirement age.

I find that the nature of social care funding plays a major role in how households choose to save for retirement. In particular, the current framework almost completely removes the need for precautionary saving for social care for households before and during retirement. My findings that household retirement planning under the current policy closely resembles what it would be if social care were provided in a universal manner, whilst representing vastly less saving than it would if social care were provide only to those with no wealth of any kind, demonstrate this. Social care funding in its current form drastically reduces the need for precautionary saving, but further generosity would be unlikely to have a similar effect.

I also find that the style of the means test used in social care funding also has a big impact on the nature of saving for retirement. I find a wide variance in the proportion of

home-owners, and the role of housing in asset portfolios of the elderly, depending on the extent to which housing wealth is included with other savings in the asset test for social care. I find that the current policy splits the difference in consequences for home ownership in the wide gap between housing always and never being included as part of savings.

3 The long-run effect of the Education Maintenance Allowance

This paper studies the short and long-run impact of a conditional cash transfer in England, known as the Education Maintenance Allowance (EMA), given to 16-19 year olds from low-income backgrounds if they stayed in full-time education, using administrative data and a reverse difference-in-differences approach. In the short run, we find that the EMA increased education participation for eligible 16-19 year olds by 2 percentage points, and reduced the labour supply of those in school. These effects are smaller than suggested by previous studies, highlighting that pilot studies can overstate true policy effects. The effects we observe were driven by an increase in the number of students doing the lowest qualification level vocational courses, with no increase in either participation in academic courses or higher level vocational courses. Despite the increase in educational participation, there is no evidence that the conditional cash transfer had long-run benefits for labour market earnings.

3.1 Introduction

A lack of basic skills can be an important drag on labour markets and productivity: lower skilled adults earn lower salaries and are less likely to work, they are more likely to claim benefits, and they are much more likely to be involved in crime. Because of this, increasing educational attainment is often a key priority for policymakers. One set of policies targeted at doing this are cash transfers conditional on educational participation. Conceptually, monetary incentives could increase the amount of people that decide to stay in education, and relieve financial concerns for students who might work while in education.

Previous work has shown that these conditional cash transfers can significantly increase education participation and reduce school dropout rates. Programmes in Australia, New York City, and England have been shown to be effective (Dearden and Heath, 1996; Dearden et al., 2009; Riccio et al., 2013). In developing economies, particularly in Latin America, cash transfers conditional on education have shown significant positive effects on both short-run academic performance (Attanasio et al., 2012; Barrera-Osorio et al., 2011; Galiani and McEwan, 2013; Glewwe and Kassouf, 2012), and on longer term outcomes such as university attendance (Barrera-Osorio et al., 2019) and earnings (Barham et al., 2013).

In this paper, we study whether conditional cash transfers in developed countries can have similar transformational effects. To do this, we evaluate the short and long-run impact of the Education Maintenance Allowance (EMA), a conditional cash transfer given to 16-19 year olds from low-income backgrounds in the United Kingdom if they stayed in full-time education. The EMA was first introduced in a pilot scheme in 1999, and was rolled out nationally in 2004. We exploit geographical and temporal variation in the rollout of the EMA via synthetic difference-in-differences in reverse (Arkhangelsky et al., 2021), using administrative data that links individual level education data across schools and colleges with long-run earnings data. Using this, we estimate the effect of living in an area that offered the EMA on short- and long-run education and earnings effects.

In the short run, we find a significant positive effect of 2 percentage points of this conditional cash transfer on education participation aged 16-19. This is smaller than that found by previous research on the EMA (Dearden et al., 2009), highlighting that pilot studies can overstate the effect of the full rollout of a policy. The effects are driven by an increase in the number of students doing vocational courses, with no increase in the number of students on the academic track. Meanwhile, the policy shifted young people away from a mix of work-based learning schemes, and from being not in education, employment, or training (NEET).

These short-run positive effects on participation are larger for boys, for students with special educational needs, and for students with worse performance in prior national exams. Chiming with the last point, we find that the increase in vocational education consists solely of an increase in the lowest level vocational qualifications, which have the lowest admissions requirements.

We also find that students aged 16-19 significantly reduced their labour supply in response to the policy by over 10%, highlighting the role of a CCT on reducing the need to work outside of school. We do not find any pass-through from reduced labour supply to improved academic performance.

In the long run, we estimate the effects of the conditional cash transfer on higher education and earnings. Despite the positive effect on education participation, and reduction in labour supply while in school, we see no effect on academic performance and no effect on university attendance. With respect to earnings over the course of the following decade, we find no significant positive effects, with generally negative point estimates, suggesting that the courses young people moved into may not have been more valuable than what they would have done without the EMA.

This paper is the first to examine the causal impact of cash transfers conditional on education participation on long-run outcomes in a developed economy. Previous studies have shown positive short-run effects of CCTs for participation in full-time education in England (Dearden et al., 2009), New York (Riccio et al., 2013), and Australia (Dearden and Heath, 1996). Other programmes focusing on training for those not in education, employment or training (NEET) have found similar positive effects in the short run (Filippucci, 2022). However, the evidence on whether these positive effects persist over the life cycle is very limited. We are the first to show that short-run increases in participation can fail to translate to long-run increases in human capital. This extends into individuals' late twenties, when we would expect additional effects even from higher education to have led to increases in earnings.

There is a wider literature on the role of conditional cash transfers in both the short and long run in developing economies (García and Saavedra, 2017). Positive impacts on short-run participation have been found in Sub-Saharan Africa (Baird et al., 2011), Asia (Barrera-Osorio and Filmer, 2016; Filmer and Schady, 2011) , and Latin America (Attanasio et al., 2012; Barrera-Osorio et al., 2011; Galiani and McEwan, 2013; Glewwe and Kassouf, 2012). In the long run, evidence from Latin America has found some significant effects on university attendance (Barrera-Osorio et al., 2019), sectoral employment shifts (Behrman et al., 2011), and earnings (Barham et al., 2013), though (Molina-Millan et al., 2016) argue that the evidence for positive effects on most long-run outcomes is inconclusive.

There is clear reason to think that conditional cash transfers would have a different effect depending on the wealth of the country they were implemented in: in wealthier countries, the budget constraint that might force students into work rather than full-time education is less likely to bind than in poorer countries. Even in low and middle income countries, there is within-country variation in the effectiveness of CCTs by income (Galiani and McEwan, 2013). Meanwhile, in England, the proportion of low-income students right after school leaving age who were not in education, employment or training was almost three times the number of students who were in full-time work, suggesting that the financial constraints may not be a key barrier to education participation.

This paper's joint analysis of the short and long-run consequences of the CCT add to the literature on education interventions with disparate short and long-run effects. There have been cases in the past where programmes with limited visible short-run impact have had positive long-run consequences, in both developing and developed countries (Baird et al., 2016; Heckman et al., 2010). On the other hand, there is also an extensive literature on interventions with positive short run impacts that have not translated into long-run benefits (Alzúa et al., 2016; Evans and Ngatia, 2021; Schochet, 2008). This paper adds to this literature, highlighting the lack of long-term effect arising from a short-run boosting of education participation.

The rest of this paper starts with an outline of the institutional background of the Education Maintenance Allowance and the post-16 education landscape in England. Sections 3.3 and 3.4 respectively describe the data and the methodology we use to estimate the effect of the EMA on a range of outcomes. Section 3.5 reports both the short and long-run impacts of the policy, and analyses the heterogeneity in effects across groups. Section 3.6 concludes.

3.2 Institutional Background

Education Maintenance Allowance The Education Maintenance Allowance (EMA) was a scheme that gave 16 to 19 year old students from low-income backgrounds a weekly stipend if they remained in full-time education. At the time of its first introduction in 1999, the compulsory school leaving age in England was 16, and so students faced a choice as to whether or not to remain in education beyond this point. It could be claimed for up to two years of education on either vocational or academic education, though not university education.

First announced in the spring of 1999, the EMA was rolled out as a set of pilot schemes in 15 out of 144 English Local Education Authorities (LEAs) in September 1999, the start of the 1999/2000 academic year. Children in these LEAs were eligible for up to £30 or £40 per week, depending on which pilot scheme their LEA was assigned and their parents' income, if they stayed in full time education. Also in September 1999, 11 LEAs were designated as control LEAs, chosen to have similar characteristics to 10 of the LEAs in the pilot. Much like the pilot areas, these were in deprived, mostly urban areas outside of London. In the next academic year, the EMA pilots were expanded to another 36 LEAs in the most deprived areas of the country, though this did not include the 11 control LEAs. For the four years after that, LEAs containing about half of the population of England, mostly in more urban areas with higher levels of deprivation, had access to the EMA, and the other half of the population, mostly in rural and suburban areas with lower levels of deprivation, did not.

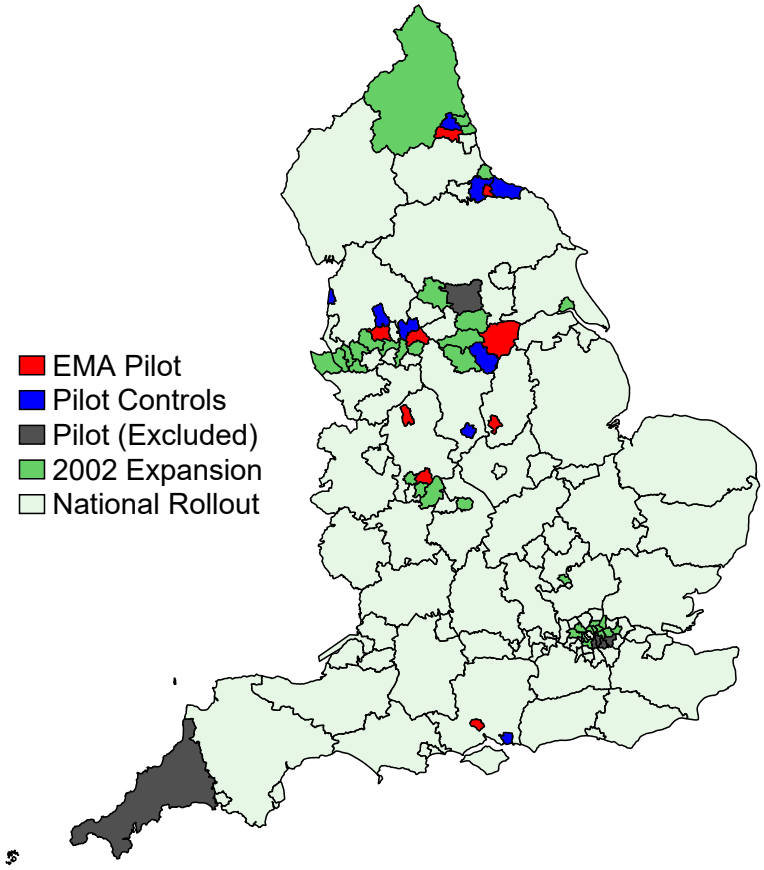

Figure 3.1 Education Maintenance Allowance Rollout

Past work on the EMA (Dearden et al., 2009; Sabates and Feinstein, 2008) exploited differences between the 1999 pilot and control areas to estimate the effect of the EMA on short term outcomes. In the main section of their paper, Dearden et al. (2009) use only the 9 urban non-London pilot LEAs and the 9 urban non-London control LEAs, due to the greater comparability of the pilot and control areas in urban areas. We follow them, and distinguish throughout between these two sets of 9 LEAs, and the other pilot and control LEAs.

Figure 3.1 shows the rollout of the EMA across areas. It demonstrates the geographical similarity between the 9 pilot and control areas, as well as the difference between the areas between LEAs that received the EMA in 2000 and earlier compared to those that received it in the national rollout.

For those whose biological parents earned a total of less than £19,630 per year (approximately 40% of the population), children were eligible for £30 per week during term time, which amounted to £1,200 per year. This represented a sizeable sum; £30 per week was equal to a quarter of the earnings of a 16-year-old working full-time on the minimum wage, and was about a third of total per-pupil spending on 16-18 education. It also represented

a doubling of the average income for students, a large increase. An additional 10% of the population, in households earning up to £24,030 per year, received £20 per week, while a final 20% whose parents earned up to £30,000 received £10 per week. They also received a £50 bonus for each school term of school they finished, as well as £50 bonus for good performance in exams.

3.2.1 Post-16 Education in England

At the time of the rollout of the EMA, those leaving school at age 16 faced a range of options. They could stay in a school setting, and do A-Levels, the national academic exams taken at age 17 and 18, which are required for university. This usually required getting 5 grades above a C at GCSE, the national exams taken at age 16, with some exceptions.

They could also attend a Further Education college, in which they could do a range of different academic and vocational courses. This included doing A-Levels, with similar restrictions as when taken in school. It also included combining some academic work with vocational qualifications, or just doing vocational qualifications. Within vocational qualifications, there were multiple levels: level 3 vocational qualifications, which are roughly equivalent to an A-Level, level 2 vocational qualifications, which are roughly equivalent to getting 5 grades above C at GCSE, and level 1 qualifications, which are equivalent to worse GCSE results, and would likely only be done by those who previously had very poor GCSE performance.

School leavers could also work, though it was very uncommon for students to do this full time immediately after leaving school. More common was doing an apprenticeship, which combined some paid work with some vocational training in an FE college. Since this education was not full-time, those doing apprenticeships were not eligible for the EMA. A final course taken by school leavers was to become NEET, and not enter education, employment or training. Some of these might return to a Further Education College in the next academic year.

Within this framework, those who remained in school or attended a further education college, except for as part of an apprenticeship or other work-based learning, were eligible for the EMA for two years of study, the standard length of time to complete A-Levels. The barriers to application were very low, and uptake was high, so we expect that almost all of those who were eligible for the EMA and in education received it.

3.3 Data

We use two different administrative datasets to estimate the effect of the EMA in both the short run and the long run. In order to estimate the effect on outcomes in education and the labour market, we use the Longitudinal Educational Outcomes dataset, that links school records on the universe of state-educated children in England with separate administrative datasets from Further Education colleges and from universities. Finally, it is linked to administrative tax and benefit data.

School records from the National Pupil Database (NPD) include data on results at nationally taken academic exams. These include SATs, taken at age 11 by all state-educated students, GCSEs, taken at age 16 by all students, and A-Levels, taken at age 18 by some students. The NPD incorporates data on whether students remained in school between after age 16. It also includes some demographic information on students. This includes month and year of birth, gender, ethnicity, first language, and special educational needs status.

The NPD also includes an indicator of whether a student is eligible for free school meals. Eligibility for this is determined by household income, and is generally limited to the bottom 15% of the population by household income. The school records do not have any marker of eligibility for the EMA, so we use eligibility for free school meals as a proxy. Using the Family Resources Survey, we estimate that over 99% of those who were eligible for free school meals would also have been eligible for the full EMA award in the same year.

Since data on free school meals comes from school records, we only have data up to the last year of compulsory schooling, age 16. We restrict our data to those who were eligible for free school meals at age 16. This means that some would not have been eligible for the EMA at ages 17 to 19, if their parents had a sudden large increase in their income, but we expect this to be rare. Therefore, we expect that when we restrict our sample to those who were eligible for free school meals, we capture a group who were almost all eligible for the full EMA, and were mostly at the lower end of the income distribution of those eligible. The results of this paper should be interpreted as the effect of the EMA on those from the bottom 15% of the household income distribution.

Data on further education comes from the Individual Learner Record (ILR). This has data on whether a student was attending a Further Education college in each academic year. It also provides data on all of the courses they studied, from which we can derive whether they were vocational or academic, and whether they completed them, giving us the level of the qualification achieved. It also has data on apprenticeships, allowing us to analyse them as a separate outcome, and exclude them from our marker of full-time education.

Data on university education comes from the Higher Education Statistics Authority (HESA). It includes data on what university students attended, what degree classification they achieved, and at what age they attended.

Finally, administrative data from HMRC, the tax authority, enables us to link educational data with longer run earnings and benefit receipt. Tax data from HMRC reports total earnings in each tax year for all individuals, allowing us to develop a picture of the earnings of students during and after education.[1] It does not include data on other aspects of employment, such as hours or occupation, so we only use annual earnings as an outcome variable.

The data from HMRC also includes information on individual benefit spells. This includes information on what benefit was received, for how long it was received, and whether it was an out-of-work benefit. This allows us to estimate the effect of the EMA on different types of benefit, as well as giving us an additional indicator of unemployment, in addition to earnings data.

The earliest cohort for whom all of this data exists are those who finished their last compulsory year of schooling in the summer of 2002, when they were 16, and so were born during the 1985-1986 school year. We use data on all cohorts up to those born during the 1994-1995 school year. Our data extends to 2020 for adult outcomes. This means that we have data on school outcomes for all individuals, as well as details on university and further education courses occurring up to age 23 for all individuals. For all cohorts we consider, we have earnings data up to age 24, though we also use earnings at older ages for earlier cohorts, up to age 28.

This data allows us to demonstrate the disparities between the different areas we consider in our evaluation of the Education Maintenance Allowance. Table 3.1 shows the variation in demographics across areas that received the EMA at different times. It demonstrates that the areas earmarked for the 1999 pilot, and the 9 control areas chosen to resemble them, looked very similar, at least from the perspective of those in the population receiving free school meals. A similar proportion of students were white, and a similar proportion did not speak English as a first language (EAL). In terms of academic achievement prior to leaving school they were also similar, with similar proportions receiving at least 5 C grades at GCSE.

There are much larger differences when comparing the areas that received the EMA before 2004 and those who got it as part of the national rollout. The one third of all LEAs that provided the EMA before 2004 had over half of the total population on free school meals.

[1] Because the data cerives from administrative tax data, we do not see earnings that are not declared to HMRC. In addition to illegal activity, this could include those who earn below the lower earnings limit (£4,004 in 2004), below which employers were not obliged to report salaries in the 2000s. However, we plot the distribution of annual earnings reported to HMRC, and we see no discontinuity at the lower earnings limit, and so we conclude that almost all employers did report earnings even below that point.

	1999 Pilot	Pilot Control	Pre 04 Rollout	2004 Rollout	All
% White	72.8	80.1	58.0	79.5	68.4
% 5 A*-C	24.0	24.8	30.4	27.2	28.8
% EAL	17.5	15.5	29.9	13.1	21.7
Obs.	49,983	39,235	387,117	366,428	753,545
LAs	9	9	51	93	144

Table 3.1 Background Characteristics of Students on Free School Meals, Born 1985-1995

This was partly because they were more populous LEAs, butalso because they had a higer proportion of students on free school meals, as the EMA was targeted more heavily at areas with higher levels of deprivation.

In other demographic characteristics, the areas also looked different. Areas that received the EMA earlier were much less white, and had a much higher proportion of students who did not speak English as a first language, both reflective of the fact that they were generally more urban. They also had slightly better prior attainment, though by nowhere near as large a margin.

Table 3.2 highlights that the same is true when considering the educational routes at age 16 across different groups. In the pilot and control groups, students had very similar probabilities of staying in full time education, doing A Levels, being NEET, and earning more than £1,000.

	1999 Pilot	Pilot Control	Pre 04 Rollout	2004 Rollout	All
% FT Education	73.2	73.4	77.0	74.4	75.8
% A Levels	28.5	27.8	36.6	31.2	33.9
% NEET	19.4	19.2	17.3	19.4	18.3
% Earnings £1k	21.5	21.5	18.9	26.3	22.5
Obs.	49,983	39,235	387,117	366,428	753,545
LAs	9	9	51	93	144

Table 3.2 Age 16 Outcomes of Students on Free School Meals, Born 1985-1995

However, the areas that became eligible for the EMA before 2004 were more likely to stay in full time education, were more likely to do A Levels, less likely to be NEET, and much less likely to be doing more than £1,000 worth of paid work. This highlights again the different landscapes that the two phases of the EMA were introduced into.

3.4 Methodology

The outcomes of interest in this paper are educational participation, qualifications, labour supply and earnings, and criminal activity. In order to estimate the effect of the Education Maintenance Allowance on these outcomes, we exploit the geographical and temporal variation in its rollout. To do this, we use a difference-in-difference methodology, revolving around the national rollout in 2004.

Unlike in a regular difference-in-difference specification, where the comparison is to units that do not receive the treatment, here we are comparing newly treated areas to areas that always received the treatment, because we do not have data on any academic cohorts that were not eligible for the EMA anywhere in England. Therefore, we use a reverse difference-in-difference specification (Kim and Lee, 2019), in which we rely on parallel trends in the outcomes we consider *post-treatment* period, rather than the pre-treatment period. The regression specification for outcome Y for individual i in cohort t is therefore:

$$Y_{i,t} = \beta_0 + \beta_1 Pre_t + \beta_2 AlwaysTreated_i + \beta_3 Pre_t \times AlwaysTreated_i + \epsilon_{i,t} \quad (3.1)$$

where Pre_t indicates that an individual is in a pre-National Rollout birth cohort and $AlwaysTreated_i$ indicates that an individual is in an area that received the EMA before the National Rollout. β_3 is the parameter identifying the Intention To Treat (ITT) effect of the EMA.

Our regression specification uses data on all students receiving free school meals in England, and using all LEAs that received the EMA before 2004 as an always-treated group, and all LEAs that received the EMA in 2004 as a switching group. The drawback of this specification, as outlined in Section 3.3, is that the areas that received the EMA before 2004 did not closely resemble those that were received in 2004. Therefore, there is a risk that trends in outcomes across the two areas may not be parallel.

Because of this, we use an adapted version of a synthetic difference-in-difference specification (Arkhangelsky et al., 2021). The synthetic difference-in-difference method involves reweighting the control group, such that for a given outcome, the treatment and control groups follow parallel trends prior to treatment. Because we use a reverse difference-in-difference, we instead reweight the switching group, such that the trends in the outcomes of interest are parallel after treatment.

This relies on a stricter set of assumptions than the standard reverse difference-in-differences. In the standard reverse difference-in-difference method, because we are using already treated areas as a control, we are reliant on there not being dynamic treatment effects

(Goodman-Bacon, 2021). If the effect of treatment in the always-treated areas were still changing upon treatment of the switching group, then we would recover biased estimates of the effect. However, this would be visible in the post-trends, as it would cause them to be non-parallel.

In the synthetic difference-in-difference case, because there already exist parallel trends for which we are not able to control, we are not able to observe whether there are dynamic treatment effects in the post-treatment trends. Therefore, we have to rely on the assumption that there are no dynamic treatment effects. Because we do not see dynamic treatment effects across the 9 pilot and control areas in the standard difference-in-difference, it seems unlikely that there would be dynamic effect across all LEAs.

Furthermore, because our analysis at the cohort level, there is little reason to think there would be dynamic treatment effects, as that would require students turning 16 different numbers of years after the implementation of the EMA to be affected differently. The main mechanism by which we would expect this to happen is through a dynamic effect on EMA uptake. Analysis of uptake statistics suggest that there was not a dynamic effect, as EMA receipt stayed relatively constant in areas once it had been rolled out.[2] Therefore, there is little reason to think there would be dynamic treatment effects that might bias our results.

As a robustness check, we also follow Dearden et al. (2009), and use a comparison between 9 of the original pilot LEAs, and the 9 control LEAs that resembled them. Because the nine control areas first received the EMA in 2004, we can explore the difference in the change in outcomes between cohorts who left compulsory schooling in 2003 and 2004 across the pilot and control areas. In the control areas, only students who turned 16 after September 2003 might have been eligible for the EMA, while in the pilot areas, all students in our sample were eligible for the EMA. This allows us to identify the change in outcomes that happened because of EMA eligibility. Because the control areas were specifically chosen to resemble the 9 pilot areas, we rely on the assumption that when all cohorts nationwide are treated from 2004 onwards, trends in outcomes are similar across the two sets of areas. We verify this assumption in the next section.

While there is a benefit to use the directly comparable LEAs of the 1999 pilot and control groups, this provides limited statistical power, due to the relatively low sample size: as Table 3.1 and 3.2 show, there were only 50,000 students on free school meals across the decade we consider in the pilot areas, and only 40,000 in the control areas, amounting to 9,000 students a year in total. This only gives us relatively imprecise estimates: for example, we can only estimate the effect of the EMA on participation in full-time education at ages 16-17 within a confidence interval with a width of approximately 6 percentage points. Therefore, we do not

[2]Statistics on uptake of the EMA can be found in Bolton (2011)

use this as our primary specification, but report the effects in Section C.1 of the Appendix. In each case, the effects reported in our main results are within the confidence intervals of the robustness checks, and in most cases they are very close.

For the standard difference-in-difference method, we cluster our standard errors at the LEA level. For the synthetic difference-in-difference method, we use the placebo method, to construct standard errors from simulated cases where a different selection of LEAs were treated.

3.5 Results

3.5.1 Effect of the EMA on education participation

Table 3.3 shows estimates of the effect of the EMA on different education routes in the academic year after compulsory schooling ended. This shows that EMA did incentivise more students to say in education beyond the compulsory schooling age. Results from the synthetic difference-in-differences suggest an increase in the number of students in full-time education of 2 percentage points. This represents a 5.6% reduction in the dropout rate at age 16 for children on free school meals.

Table 3.3 Impact of the EMA on education participation aged 16-17 (FSM students)

	Baseline	Synthetic DiD
Full Time Education	0.640	0.020***
		(0.007)
Further Education	0.476	0.032***
		(0.009)
In School	0.187	-0.002
		(0.006)
A-Levels	0.232	-0.004
		(0.006)
In Training	0.066	-0.013***
		(0.004)
NEET	0.281	-0.012*
		(0.006)
Observations		753,545
LAs $\times$ Years		1440

Age 16-17 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

This appears to be driven by an increase in the number of students pursuing degrees in Further Education colleges, where vocational courses are more common. There is no change in the number of students staying in school, nor is there any change in students remaining on the academic track (either in schools or in FE colleges). Part of the increase in education participation seems to mostly come from a significant drop in the number of students who are not in education, employment or training (NEET), the rates of which decrease by 4%. This suggests that the EMA drew some people away from less productive activities, and towards education. It also suggests that most of those who could be kept in education by the EMA were not out of education because they were financially constrained, as they were not otherwise earning money. There is also a decrease of 1.3 percentage points in the number of students in work-based training, an drop of almost 20%. Work-based learning was not covered by the EMA, and so the fall in training was likely driven by an increase in the relative financial value of full-time education relative to work-based learning.

Notably, the point estimates of these results are lower than the estimated effect of the EMA in other studies. Dearden et al. (2009) find a 6.7 percentage point increase in participation in full-time education in the first year after leaving compulsory education, amongst those eligible for the full EMA. There are several possible explanations for this. Firstly, our estimates are based just on those receiving free school meals, the lowest income group receiving the EMA, whereas their study incorporates all students eligible for the full EMA. It may therefore be the case that effects are larger for higher income students with full eligibility. However, this is unlikely, as we should expect students with lower incomes to be, if anything, more responsive to income incentives.

There are several other explanations: firstly, the effect may have been larger in the original pilot areas, which they use; indeed, the standard method using the pilot areas in Table C.1 in the Appendix shows a larger effect of 0.036 percentage points. Given the relatively small population underlying both Dearden et al. (2009) and the difference-in-difference method using the pilot and control areas, the confidence intervals are wide, and the difference between these results is not significantly different. Finally, some of the variation may be explained by changing effect over time. The value of the EMA did not increase with inflation, so became less valuable over time: this might explain a difference in effect between 1999 and 2004. It would also imply the effect of the EMA might have further reduced going forward in time, as the nominal value was never increased in England.

Table 3.4 shows that the effect of the EMA on education choices in the second year after the end of compulsory schooling was very similar. Participation in full-time education increased by 1.7 percentage points, representing a reduction in the dropout rate from 16 to 18 of 3.4%. Again, these results are concentrated in an increase in attendance of FE colleges,

Table 3.4 Impact of the EMA on education participation aged 17-18 (FSM students)

	Baseline	Synthetic DiD
Full Time Education	0.499	0.017**
		(0.007)
Further Education	0.397	0.022***
		(0.008)
In School	0.115	-0.005
		(0.005)
A-Levels	0.209	-0.003
		(0.006)
In Training	0.091	-0.014***
		(0.004)
NEET	0.343	-0.001
		(0.006)
Observations		753,545
LAs $\times$ Years		1440

Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

with no significant change in the number of students in schools, or doing A-Levels. Unlike the year before, we do not see a significant decrease in the number of students who are NEET. However, we still see a decline in work-based learning.

We break down these results according to different dimensions of heterogeneity, in order to see for whom the EMA was particularly effective for. Firstly, we break down the results by gender, shown in Table 3.5. In the two school years following the end of compulsory schooling, we find a larger positive effect on boys than on girls. This matches the similar findings of Dearden et al. (2009), who find significantly larger effects of the EMA on participation for boys than for girls.

Next, we examine the differential effect of the EMA by prior academic attainment. For this, we divide the population by whether they achieved 5 grades of between A* and C at GCSE, the standard threshold for continuing on to do A-Levels, and the equivalent of a Level 2 vocational qualification. The results are presented in Table 3.6. They show a large disparity in effect sizes by prior academic performance: those with better prior academic performance were not significantly affected by the EMA. Meanwhile, those with worse prior attainments experienced reasonably large effects. This is perhaps not surprising given the lack of effect on A-Levels: in areas without the EMA, those with 5 A* to C grades were 59 percentage points more likely to do A-Levels.

Table 3.5 Impact of the EMA on education participation by gender (FSM students)

	Synthetic DiD (Age 16-17)	Synthetic DiD (Age 17-18)
Boys' Participation	0.027***	0.027***
	(0.010)	(0.010)
Observations	382,605	382,605
Girls' Participation	0.018**	0.010
	(0.009)	(0.008)
Observations	370,940	370,940
LAs × Years	1440	1440

Age 16-17 and Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

It is worth noting that participation rates differ significantly with academic attainment. Only 60% of those who did not have 5 C or better grades at GCSE remained in full-time education, whereas 83% of those who did did. This means that there was much more room to increase participation for those with lower qualifications. However, we should expect returns to education to be higher for those with higher previous attainment, and so it is noteworthy that those who could benefit most from additional education were not the most likely to be induced into education participation by the EMA.

Table 3.6 Impact of the EMA on education participation by prior academic performance (FSM students)

	Synthetic DiD (Age 16-17)	Synthetic DiD (Age 17-18)
5 A*-C Participation	0.009	0.016
	(0.007)	(0.010)
Observations	217,375	217,375
Not 5 A*-C Participation	0.031***	0.024***
	(0.009)	(0.009)
Observations	536,170	536,170
LAs × Years	1440	1440

Age 16-17 and Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

As a final piece of heterogeneity analysis, we analyse the difference between children receiving support for Special Educational Needs (SEN), and those not. This is a marker of those diagnosed with learning difficulties. Among others, these include dyslexia, attention deficit hyperactivity disorder, autism, and physical disabilities. Table 3.7 shows the full breakdown of effects.

The effect of the EMA was much larger for those receiving support for special educational needs. As with the difference in effect by attainment, this might derive partly from the

different baselines: education participation the year after compulsory schooling for those with SEN was 59%, while participation for those without SEN was 68%.

Table 3.7 Impact of the EMA on education participation by special educational needs (FSM students)

	Synthetic DiD (Age 16-17)	Synthetic DiD (Age 17-18)
SEN Participation	0.053***	0.040***
	(0.013)	(0.013)
Observations	260,856	260,856
No SEN Participation	0.013*	0.015**
	(0.007)	(0.007)
Observations	492,689	492,689
LAs × Years	1440	1440

Age 16-17 and Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

3.5.2 Effect on short-run labour supply

We should not only expect the EMA to have an effect through the complier population, who only remained in full-time education because of the EMA, and would otherwise not have stayed in school. There is also reason to believe that the "always-takers" of education would also be affected, through an income effect on labour supply.

We do not have good data on labour supply in the form of hours or days worked, so we use annual earnings as a proxy. Given they all have very similar qualifications, and are almost all working part-time, this is likely to be a good approximation of labour supply, due to limited variation in hourly wages. Additionally, it is unlikely that the EMA would have a significant effect on hourly wages in the short run, so we should expect increased earnings to Table 3.8 shows the effect of the EMA on earnings at age 17. Since the tax year, the unit of time for which we have earnings data, does not align with the academic year, this reflects earnings across the part of two academic years, including the last five months of the first post-compulsory academic year, and the first seven months of the second. At some point in this period, all individuals in the data are 17.

The results show that there is a significant and large drop in labour supply in response to the EMA. The average 17-year-old earned £145, or 8%, less as a result of the EMA. It is possible that this might just be driven by an increase in the number of students, and thus a decrease in the number of students in work-based learning and full-time work. We see a reduction of 1.3 percentage points in the number of students earning above the lower earnings limit. In 2004, the first year of the EMA, this was £4,004. A 17 year old earning

the minimum wage for their age would have to work 25 hours a week to earn above the threshold, and so it is unlikely that they were previously working part time. Therefore, this likely reflects the fall in work-based learning.

However, we can also see that the reduction in total labour supply does not just derive from the extensive margin. The effect of the EMA on the earnings of those who remain in full-time education is £153, a reduction of 12%, larger than the effect on the population as a whole. This reflects that those who would have stayed in education even if the EMA had not existed significantly reduced the amount of part time work they did because of the extra £30 per week of increased income. The additional £30 per week was equivalent to £1,200 per year, meaning that the labour supply response to the income shock amount amount to 15% of the total income shock. It does not appear that any of this adjustment comes from those who would otherwise have done some time part-time work shifting to doing none, as exemplified by the fact there was no effect on the number of students reporting any earnings at all.

Table 3.8 Impact of the EMA on work at age 17 (FSM students)

	Baseline	Synthetic DiD
Annual Earnings, Age 17, £	1842	-145***
		(54)
% Earning above LEL, Age 17	0.135	-0.013***
		(0.004)
Earnings above 0, Age 17	0.417	-0.005
		(0.008)
Observations		753,545
Annual Earnings, Age 17, £ In Education	1320	-153*** (50)
Observations		570,700
LAs × Years		1440

Age 17 represents the tax year (April-April) in which everyone in the cohort is aged 17 at some point. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

3.5.3 Effect of the EMA on qualifications

Given the increase in participation and the decrease in labour supply, there are several mechanisms by which we might expect the EMA to positively affect qualifications. We consider three different forms of qualification. Firstly, we estimate the effect of access to the EMA on performance in A-Level exams, for those who take them. This might be where one would expect to see the effect of reduced labour supply in improved grades. However, as

Table 3.9 shows, the EMA did not have a signifcant positive effect on A-Level results. If anything, the effects are very slightly negative, though only significant at the 10% level.

Next, we examine the effect on the number of students achieving different levels of technical qualifications. There are three levels which are commonly achieved by age 18: Levels 1, 2 and 3. Level 2 is equivalent to getting over 5 Cs at GCSE, and Level 3 qualifications are equivalent to A-Levels. We break down whether the EMA increased the probability of students getting above each level of qualification. This shows that the only effect is on Level 1 qualifications, which increased significantly in response to the EMA, roughly in line with the total increase in full-time education. There is no significant effect on the number of students achieving a Level 2 qualification or above.

This suggests that the increase in educational participation as a result of the conditional cash transfer is driven exclusively by students entering the lowest level of qualifications. These qualifications are reasonably common amongst students on Free School Meals, being done by just over a third of students. Given the results of Table 3.6 that the increase in participation occurs only amongst those without 5 A*-C grades at GCSE, it is not surprising that those who enter into education because of the EMA generally do less advanced qualifications.

Finally, we estimate the effect of the EMA on university attendance. The positive effects on participation in education from 16 to 18 do not translate into greater participation in higher education. This is a consequence of the fact that the increase in participation is located in Level 1 courses, which are not sufficient qualifications to attend university. When combined with the lack of improvement in academic performance amongst those doing A-Level qualifications, all clear mechanisms by which the EMA might have improved university attendance are ruled out.

3.5.4 Effect of the EMA on long-run earnings

To estimate how the short run effects of the Education Maintenance Allowance translated into long-run human capital, we estimate the effect of exposure to the EMA on earnings at a range of ages. The first effect we estimate is of exposure to the EMA on the logarithm of earnings at ages 22, 24, 26, and 28. The results are shown in Table 3.10. In each case, this reflects the effect of the EMA on annual earnings, with those reporting no earnings removed. We observe no significant effect on earnings at any age, with most point estimates negative. This suggests that the long-run impact of the EMA on wages was limited.

Given that the conditional cash transfer almost exclusively brought students into the lowest difficulty qualifications, the lack of increased earnings in response might partly reflect limited returns to these courses. This is particularly true given the fact that many of these students entered education instead of doing apprenticeships, which one might think would

Table 3.9 Impact of the EMA on qualifications (FSM students)

	Baseline	Synthetic DiD
A Level Results, Standardised	-0.945	-0.062*
		(0.034)
Observations		256,010
Level 1 Qualification Or Above	0.691	0.022**
		(0.009)
Level 2 Qualification Or Above	0.328	0.008
		(0.007)
Level 3 Qualification Or Above	0.142	-0.004
		(0.005)
University Attendance	0.102	-0.008
		(0.006)
Observations		753,545
LAs $\times$ Years		1440

Level qualifications considered by age 19, and only include those obtained at FE colleges, not schools. University degrees considered by age 23. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

have relatively high returns. However, given the precision of the results in Table 3.10, only limited conclusions can be drawn about the return to Level 1 Qualifications. At the 5% significance level, only effect sizes larger than a 2% increase in earnings at age 26 would be visible. The "compliers", those who only stayed in education because of the EMA, make up only 2% of the sample. If increases in earnings only came from the compliers, the returns to level 1 qualifications would need to be over 100% for the effect of this additional education on their earnings to be visible in our results.

Nevertheless, there are other ways in which one might the the EMA could have affected earnings in the long run. Reduced labour supply, or other stresses alleviated by additional income, might have improved performance, with long-run human capital gains, that might not be visible in exam results. In spite of this, it appears there was no effect. To show this in another way, we present the effect of the EMA on earning over a rage of thresholds at ages 17, 22, and 26 in Figure 3.2.

The leftmost chart shows again that earnings over the lower earnings limit and over £10,000 are reduced at age 17, a result likely driven by increased participation and reduced work-based learning. However, at age 22 we see no significant positive effects, and see a significant negative effect on earning over £20,000. Meanwhile, at age 26 we also see no significant effects, but the point estimates for earning above each value are negative in each case. From this, it does not appear that there is a positive effect on earnings from the EMA, and that there may even be a negative effect at certain points in the distribution.

Table 3.10 Impact of the EMA on earnings (FSM students)

	Baseline	Synthetic DiD
	Age 22	
Earnings (Log)	£9200	0.007
		(0.013)
	Age 24	
Earnings (Log)	£10,500	-0.016
		(0.013)
	Age 26	
Earnings (Log)	£11,400	-0.007
		(0.010)
	Age 28	
Earnings (Log)	£12,700	0.001
		(0.013)

Results for those who have annual earnings greater than zero in the given year. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

There are several possible explanations for this phenomenon. The most straightforward of these lies in the fact that many of those who entered into education because of the EMA did so as a replacement for work-based learning, which might have had higher returns to future earnings than the lower level qualifications students did instead. There may also be other explanations, such as peer effects: the presence of more students with low prior attainment might have disrupted learning environment for others, or higher attaining students might have made similar educational choices to their friends, who otherwise would have left education.

3.6 Conclusion

Previous work has shown the effectiveness of conditional cash transfers in incentivising education participation across a range of concepts, including that of the Education Maintenance Allowance. This paper is the first in a developed country to study whether this short-run effect translates to long-run benefits to students. Despite finding modest but positive short-run effects, this paper finds no evidence that this led to long-run benefits to human capital and wages.

Overall, the short-run effects are positive: the EMA did increase participation in education, reducing dropout rates at age 16 by almost 6%, though this effect is smaller than that found by previous work by Dearden et al. (2009). Furthermore, for those already in school the

Figure 3.2 Effect on Proportion Earning More Than A Certain Level At Different Ages

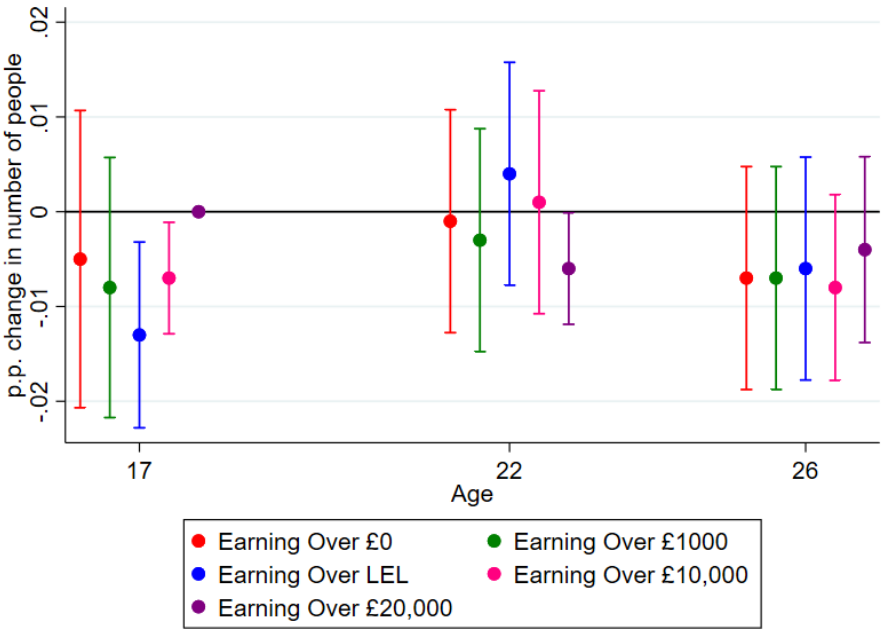

conditional cash transfer also made a difference, reducing labour supply for those in education by more than 10%.

However, the top-line results of the 2 percentage point increase in education participation as a result of the EMA are concentrated amongst specific groups, and confined to specific categories of education. Groups with lower prior attainment were much more likely to be brought into education by the EMA, and only courses of the lowest difficulty level saw increases in enrollment. Given the limited long-run benefit to students from these increases, this highlights the importance of policy to encourage students into the most beneficial types of education, and not just any education. Meanwhile, the reduced labour supply while in school did not, on average, translate to improved academic performance.

Our results highlight another reason why the EMA might not have been effective in improving long-run outcomes. Those who entered full-time education in response to the CCT generally substituted away from one of two things: being in training, or being in neither employment, education or training. Those in work-based training may have been taking away from a path with higher returns by the design of the policy to only allow for students in full-time education. [3] Meanwhile, the large reduction in the number of students who were NEET suggests that without the EMA, those who were brought into education were not

[3]However, there is also no evidence to suggest that without this change, we would have seen any increase in earnings, as subgroups of the population that did not see a drop in work-based learning also saw no long-run increase in earnings.

credit-constrained, and were not in education for other reasons. This reduces the likelihood that they could expect major returns from staying in education, as they would likely have stayed in education anyway if this were the case.

Concluding Remarks

The key overarching theme of this thesis is the long-run dynamic consequences of government policy; from childcare to adult social care to support for young people, many of the key consequences of the policy occur decades away from when the policy is actually delivered. In the cases of childcare and support for young people, this is in the decades after, while in the case of adult social care, we should expect to see consequences of the policy in the decades before an individual actually needs it.

The first chapter considers the broad dynamic implications of a childcare cost subsidy in the UK, both in its effects on short-run labour supply decisions and on long-run labour market outcomes. The life-cycle model used allows for the estimation of unobservable variables regarding the utility cost of work, the accumulation of human capital, and access to non-market childcare, each of which is key for determining the role childcare costs play in household decisions. Obtaining an estimate of all three of these processes allows me to use the life cycle model to estimate the dynamic effects of a childcare subsidy.

I find that the subsidy does have a significant impact on labour market outcomes, through the mechanism of making it easier for mothers who do not have access to any form of other childcare besides market-rate childcare to work. This creates a short-run increase in female labour supply for women, and this increased labour market participation has some pass through onto long-run wages. The chapter also shows that a childcare subsidy could drive a large reduction in the child penalty in earnings, which I measure using a comparison between mothers and non-mothers of the same age. This is caused both by a reduction in the gap in total earnings due to an increase in total hours worked among mothers, but also by a long-run drop in the wage gap between the two groups, due to increased human capital accumulation among mothers.

In the second chapter, I consider the dynamic implications of social care costs, both in their broad influence on saving and in the specifics of the English social care funding context. The chapter estimates the process of need for social care and uses Simulated Method of Moments to estimate household preferences for different saving drivers and for social care

itself. This allows me to estimate the interaction between social care risks and saving choices for those of retirement age over a very long timespan.

The key finding is that social care funding plays a major role in how households choose to save for retirement - in particular, I find that the current framework of social care in England almost completely removes the need for precautionary saving for social care for households before and during retirement. The nature of the means test used in social care funding also has a big impact on the nature of saving for retirement. In particular, the current system heavily incentivises households to retain housing wealth, sometimes at the expense of extra consumption, and we should expect social care reforms that do not give special treatment to housing wealth to change the assets in which retiring households choose to store their wealth.

In the third chapter, we expand the literature on conditional cash transfers in developed economies to the long term - previous work has shown the effectiveness of conditional cash transfers in incentivising education participation, and we examine how this translates to productivity over the long run. Despite finding modest but positive short-run effects, the chapter finds no evidence that this led to long-run benefits to human capital and productivity.

The increase in education participation we see as a result of the EMA are concentrated amongst specific groups, and confined to specific categories of education. Groups with lower prior attainment were much more likely to be brought into education by the EMA, and only courses of the lowest difficulty level saw increases in enrollment. Given the limited long-run benefit to students from these increases, this highlights the importance of policy to encourage students into the most beneficial types of education, and not just any education. Furthermore, it highlights the importance of unintended consequences of policies - the cash transfer was only given to those in full-time education, which disincentivised young people from doing apprenticeships, which often had very high returns.

Overall, one of the key themes of this thesis is the necessity of evaluating policy through a long-termist lens. Across this thesis, some of the work explicitly evaluates a policy long after it has occurred. However, it is also important to think about ways to evaluate long-term consequences of policy using only information gathered over a short timespan. Without this, we would not be able to evaluate the key outcomes of many policies on a practical timeline. Structural life-cycle modelling can, in many cases, provide support here: the first two chapters of this thesis use data across a wide range of ages, and short longitudinal elements, to estimate what changes to outcomes at younger ages do to later life outcomes. However, the third chapter shows that signs as to what we might expect in the long run can be seen in many short run outcomes - in that example, increases in full time education at low levels at the expense of work-based learning should not necessarily be seen as a success for the policy in the short run.

This work also highlights the role that good longitudinal data can play in enhancing our understanding of these issues. The third chapter of this thesis shows what can be done with large administrative datasets - large sample sizes and a lack of attrition allows us to precisely and accurately capture the effect of a policy, and this could be applied to other areas. The survey data used in the first two chapters provides valuable information, but survey data often faces challenges, including attrition and measurement error, that are particularly important when estimating changes over time. High-quality and long-running longitudinal datasets will be very valuable in furthering our understanding of how the effect of public policies develops over time.

A Appendix to: Childcare subsidies and the child penalty: a life-cycle analysis

A.1 Solution of the Model

A.1.1 Use of the Endogenous Grid Method

The estimation and simulation of the model requires the solution of women's decisions across the model. Since there is no analytical solution to the problem, I solve the problem numerically using backwards recursion, starting from the end of life at age 75.

For efficiency, I use the endogenous grid method to estimate saving and consumption decisions, following Carroll (2006). However, the model involves simultaneously solving for consumption and labour supply choices, with discrete labour supply and continuous consumption choices. The use of endogenous grid method to solve models with both discrete and continuous choices is much harder, due to discrete choices creating kinks in the value function. First, it requires calculating the Euler equation for every possible discrete labour supply choice in the first instance. This removes the risk of kinks appearing in the same period as the simultaneous choice, as optimal saving decisions are made for each possible labour supply choice.

However, in this setup, you might also expect kinks in the marginal utility of consumption in future periods. The reason they might occur is that a change in present-day consumption might change future labour supply decision, which would in turn change the marginal utility of consumption in the opposite direction. This means the marginal utility of wealth is not smooth, which means that attempting to solve the model using Euler equations creates problems. To get around this, I follow the solution proposed by Fella (2014). This involves finding the bounds of the assets space in which the consumption choice involves kinks in the marginal utility of wealth, and removing the non-optimal solutions within those bounds. In this way, I obtain a solution to the model.

A.1.2 Details of Solution and Simulation

I obtain the solution by using a discrete grid to represent the state space, across which I then interpolate upon simulating the population. I use a grid of 13 points for assets, 9 points for female income shock and 9 points for male income shock. I also use 4 points for human capital, with a grid that varies depending on the age of the individual, and the human capital parameters being estimated. I use 4 grid points for number of children, and 17 for the age of the youngest child, as well as 2 for the question of whether the individual has access to non-market childcare. Because they are mutually exclusive from other elements of the state space, I append one extra grid point on the female income shock and male income shock, which represent unemployment and being unmarried respectively.

In each period, for each plausible point in the grid, I calculate the optimal consumption choice for each possible labour supply choice.[1] This takes into account expectations for future periods, including over income, probability of marriage transition, probability of fertility, as well as knowledge of what their human capital next period will be given their labour supply decision in this period. [2] Having calculated optimal consumption decisions for each labour supply choice, I then choose the labour supply choice which maximises the value function for each given consumption choice. I then input this into the value function, and iterate backwards onto the period representing an individual being one year younger. I iterate this process starting from age 75 down to age 20 (though the labour supply decision and any shocks only begins at age 65, making the first ten-periods analytically solvable), until I have a solution grid of optimal consumption choices for each labour supply choice in the state space for each age, and the value of the value function for each of these labour supply choices at each point.

I use this set of optimal choices to simulate the model. First I generate a population of 12,000 individuals across 45 years, and apply shocks to marriage, fertility, income and unemployment to them. I then apply the solution grid to the simulations. I do this by interpolating across up to four dimensions (assets, experience, and productivity shocks for both wife and husband), imputing consumption and labour supply decisions from the shocks and state variables of the simulated individuals, by linear interpolation.

[1] Implausible grid points include those like having a 14-year-old youngest child at age 23, or having a 3-year-old child at age 60.

[2] In the case of income and human capital accumulation, this is done by interpolating across the grid-points in the next period. In the case of income, this also involves using Gauss-Hermite quadrature to estimate the distribution of expectations across normally distributed shocks.

A.2 Estimation

A.2.1 Estimation of Male Income Process

I use GMM to estimate the male income process. I do this first by regressing annual earnings $w_{i,t}$ on age and age squared, as well as controls. To enable comparison with the model, I use the age of the man's wife, and not the man himself. The resulting regression is as follows:

$$\log(w_{i,t}^M) = log(w_{initial}^M) + \gamma_1 t - \gamma_2 t^2 + \xi Year_{i,t} + u_{i,t} \tag{A.1}$$

This gives an estimate of initial income, as well as the time-varying parameters. I then use the residuals u_t of this regression to estimate the other elements of the male income process, including the variance of initial heterogeneity σ_f^2, persistence ν, variance of shocks σ_ζ^2 and measurement error σ_m^2. I do this by GMM, using the following moments:

$$E[\Delta u_t^2] = 2\sigma_m^2 + \frac{1+\nu(\nu-1)^2(1-\nu^{2(t-2)})}{1-\nu^2}\sigma_\zeta^2(\nu-1) \tag{A.2}$$

$$E[\Delta u_t \Delta u_{t-1}] = \frac{1+(1-\nu)^2(1-\nu^{2(t-1)})}{1-\nu^2}\sigma_\zeta^2 - \sigma_m^2 \tag{A.3}$$

$$E[u_t^2] = \sigma_f^2 + \sigma_m^2 + \frac{1-\nu^{2t}}{1-\nu^2}\sigma_\zeta^2 \tag{A.4}$$

$$E[u_t u_{t-1}] = \sigma_f^2 + \frac{1-\nu^{2t}}{1-\nu}\nu\sigma_\zeta^2 \tag{A.5}$$

This gives the following results, shown in Table A.1:

Parameter		Estimate	Standard Error
Initial log wage	$\log(w_{Initial}^M)$	9.534	0.036
Variance of wage shocks	σ_ζ^{2M}	0.039	0.008
Variance of initial heterogeneity	σ_f^{2M}	0.159	0.009
Persistence of wage process	ν^M	0.603	0.009
Variance of initial heterogeneity	σ_m^{2M}	0.172	0.0013
Returns to age	γ_1	0.028	0.003
Returns to age squared	γ_2	0.00060	0.00006

Table A.1 Male Income Process Parameters

A.2.2 Marriage Parameters

In the Understanding Society data, there is a mismatch between the proportion of women who are married in period t and not married in period $t-1$, and the number of people who are married in total in period t. This indicates that those who got married in the last year are less likely to appear in the Understanding Society data. To accurately match the proportion of people who are married at each age from the data, I therefore impute the proportion of people who transition into marriage from the total number of married people, and the number of people who transition from marriage to singlehood in the data. This gives the following transition matrix for marriage shown in Table A.3. It represents the probability for a woman in the given age group transitioning in any given year.

Age group	$P(Marriage)$	$P(Divorce)$
20-24	0.160	0.159
25-29	0.150	0.054
30-34	0.084	0.030
35-39	0.062	0.025
40-44	0.044	0.021
45-49	0.021	0.015
50-54	0.025	0.010
55-59	0.033	0.014
60-64	0.019	0.009

Table A.2 Marriage Probability Parameters

A.2.3 Fertility Probabilities

A similar issue to that of marriage transitions exists in the data on having children in the Understanding Society data. This is exemplified by the fact that there are a much larger number of women with 1-year-old children than women with 0-year-old children in the data. Therefore, I adjust the fertility probabilities so as to match the proportions of women at each age who have 1, 2, or 3 children, for each marital status. This gives the following probabilities, shown in **??**, of childbirth in any given age group for each year, for a given marital status and number of children pre-fertility.

Age Group	Married			Single		
	Number of Children					
	0	1	2	0	1	2
20-24	0.215	0.370	0.212	0.087	0.168	0.201
25-29	0.189	0.267	0.105	0.047	0.030	0.045
30-34	0.207	0.175	0.010	0.031	0.045	0.005
35-39	0.079	0.091	0.034	0.007	0.018	0.053

Table A.3 Probability of Childbirth

A.3 Model Fit

A.3.1 List of Moments

In this section I detail the fit of the model, and the specification of the moments used. I outline the moments and their confidence intervals in the data, and the equivalent moment in the simulated data. In each case the confidence intervals come from bootstrapping the moments in the data.

The first moments are those based around the employment rate of women with different family backgrounds. This is dependent on marital status, number of children and age of children. So as to control for selection effects, including age, I use a probit model to estimate this. The specification of the probit is as follows:

$$Employment_{i,t} = \beta_1 Type_{i,t} + \beta_2 Age_{i,t} + \beta_3 Age^2_{i,t} + \beta_4 X_{i,t} + \epsilon_{i,t} \tag{A.6}$$

where $Employment$ is measured at the extensive margin, $Type$ is a dummy variable representing category of family demographics, and X is a set of controls, including year and ethnicity. After running this probit, I convert it to marginal effects, so the results represent the change in probability from changing type. The results of the value of β_i in the marginal effects regression are reported in Table A.4 below. Note that the case of being single and having no children is used as a baseline, so no coefficient is reported.

I also use other moments to capture parameters involved in the estimation of the utility cost of work. The first of these concerns single childless women, who are used as a baseline above. Therefore, I include a moment that is the mean employment of single childless women in their late twenties. In order to estimate the utility cost of part-time work for married and single women, I also estimate the ratio of full-time to part-time workers by marital status, as well as grouping by whether women are above 40 or below 40. These are shown below in Table A.5:

Type	Number	in Data	in Model	Number	in Data	in Model
		Married			Single	
No children	M_1	-0.005	-0.006		Baseline Case	
		(-0.021,0.011)				
1 child (Under 5)	M_2	-0.091	-0.084	M_{12}	-0.327	-0.329
		(-0.109,-0.073)			(-0.349,-0.305)	
1 child (5-10)	M_3	-0.010	-0.036	M_{13}	-0.195	-0.195
		(-0.034,0.014)			(-0.228,-0.162)	
1 child (10-15)	M_4	-0.054	-0.056	M_{14}	-0.088	-0.096
		(-0.072,-0.036)			(-0.108,-0.068)	
2 children (Under 5)	M_5	-0.241	-0.231	M_{15}	-0.543	-0.541
		(-0.257,-0.225)			(-0.574,-0.512)	
2 children (5-10)	M_6	-0.089	-0.098	M_{16}	-0.287	-0.301
		(-0.109,-0.069)			(-0.322,-0.252)	
2 children (10-15)	M_7	-0.033	-0.041	M_{17}	-0.088	-0.082
		(-0.055,-0.011)			(-0.123,-0.053)	
3 children (Under 5)	M_8	-0.437	-0.424	M_{18}	-0.590	-0.592
		(-0.464,-0.410)			(-0.627,-0.553)	
3 children (5-10)	M_9	-0.260	-0.294	M_{19}	-0.270	-0.279
		(-0.295,-0.225)			(-0.344,-0.196)	
3 children (10-15)	M_{10}	-0.343	-0.478	M_{20}	-0.136	-0.200
		(-0.419,-0.267)			(-0.283,0.011)	
Adult Children	M_{11}	-0.020	-0.014	M_{21}	0.000	0.010
		(-0.034,-0.006)			(-0.014,0.014)	

Table A.4 Employment Coefficient by Family Demographics

Confident intervals listed under data estimates in brackets.

In order to estimate the human capital accumulation moments, I use two panel data regressions, along with other moments. The two regressions are as follows:

$$log(w_{i,t}) = \tau_1 Experience_{i,t} + \tau_2 Experience_{i,t}^2 + \tau_3 X_{i,t} + \gamma_i + \epsilon_{i,t} \tag{A.7}$$

$$log(w_{i,t} - w_{i,t-1}) = \chi_1 FullTime_{i,t-1} + \chi_2 log(w_{i,t-1}) + \chi_3 Age_{i,t} + \chi_4 Age_{i,t}^2 + \chi_5 X_{i,t} + \epsilon_{i,t} \tag{A.8}$$

The first of these regresses experience on log wages. For experience, I use the total number of years spent working. The second regresses the type of work done last period, and whether it is part-time or full-time, on wage growth, controlling for previous wages and other observables. The baseline case is working part-time. The fit of these regressions is shown in Table A.6.

In order to estimate the accumulation of experience and the parameters of the income process, I also estimate the mean and standard deviation of log wages at five year intervals throughout the life-cycle. The fit of the model to these results is shown in Table A.7 below:

Moment	Number	in Data	in Model
Mean Employment, Single Childless Women	M_{22}	0.842	0.871
		(0.817,0.877)	
Full-Part Time Ratio, Married Under 40	M_{23}	1.535	1.244
		(1.337,1.733)	
Full-Part Time Ratio, Married Over 40	M_{24}	1.911	1.844
		(1.711,2.111)	
Full-Part Time Ratio, Single Under 40	M_{25}	2.597	2.148
		(2.158,3.036)	
Full-Part Time Ratio,Single Over 40	M_{26}	3.351	3.843
		(2.818,3.884)	

Table A.5 Remaining Employment Moments

Confident intervals listed under data estimates in brackets.

Moment	Number	in Data	in Model
τ_1	M_{27}	0.0096	0.0117
		(-0.0004,0.0196)	
τ_2	M_{28}	-0.00021	-0.00018
		(-0.00027,-0.00015)	
χ_1	M_{29}	0.0344	0.0409
		(0.0293,0.0395)	

Table A.6 Human Capital Regressions

Confident intervals listed under data estimates in brackets.

The final moments contribute to estimating the income process, and pin down the proportion of women with access to non-market childcare. In the case of non-market childcare, it should be noted that here, working means working full-time, and using childcare means using more than 20 hours of childcare with a child under the age of 4. The results are given in Table A.8 below:

A.3.2 Relevance of Moments

I now outline the importance of different moments for determining value of each parameter. I do this by using methods pioneered by Andrews et al. (2017) and Honore et al. (2020). In particular, I measure how sensitive the standard errors of the parameter values are to removing each moment from their estimation. By measuring the change in the standard errors when a moment is removed, I gain a sense of how significant a moment is to that parameter: if the standard error of the parameter would double were the moment not there, then that moment can be considered as very important for the estimation of that parameter.

Age	Number	in Data	in Model	Number	in Data	in Model
	Mean			Standard Deviation		
20	M_{30}	9.517	9.489	M_{39}	0.212	0.214
		(9.501,9.533)			(0.192,0.232)	
25	M_{31}	9.557	9.602	M_{40}	0.212	0.259
		(9.545,9.569)			(0.200,0.224)	
30	M_{32}	9.661	9.665	M_{41}	0.275	0.273
		(9.647,9.675)			(0.265,0.285)	
35	M_{33}	9.710	9.709	M_{42}	0.306	0.289
		(9.696,9.724)			(0.296,0.316)	
40	M_{34}	9.730	9.712	M_{43}	0.330	0.302
		(9.718,9.742)			(0.320,0.340)	
45	M_{35}	9.722	9.706	M_{44}	0.320	0.308
		(9.708,9.736)			(0.314,0.326)	
50	M_{36}	9.709	9.703	M_{45}	0.313	0.314
		(9.699,9.719)			(0.303,0.323)	
55	M_{37}	9.671	9.678	M_{46}	0.299	0.318
		(9.661,9.681)			(0.287,0.311)	
60	M_{38}	9.652	9.662	M_{47}	0.298	0.311
		(9.634,9.670)			(0.284,0.312)	

Table A.7 Income Statistics

Confident intervals listed under data estimates in brackets.

Due to space, I do not report every combination of 33 parameters and 51 moments, but only those that might be intuitively significant to each parameter.

Parameter	Variance	Moment Removed	Change in Variance
Utility cost of work when:			
Married, 1 child (Under 5)	0.0016	M_2	0.0029
Married, 1 child (5-10)	0.0004	M_3	0.0043
Married, 1 child (10-15)	0.0025	M_4	0.0714
Married, 2 children (Under 5)	0.0004	M_5	0.0030
Married, 2 children (5-10)	0.0004	M_6	0.0051
Married, 2 children (10-15)	0.0001	M_7	0.0.0027
Married, 3 children (Under 5)	0.0004	M_8	0.0228
Married, 3 children (5-10)	0.0001	M_9	0.0072
Married, 3 children (10-15)	0.0009	M_{10}	0.0006
Married, no children	0.0009	M_1	0.0008
Married, adult children	0.0001	M_{11}	0.00002
Single, 1 child (Under 5)	0.009	M_{12}	0.0153

Single, 1 child (5-10)	0.0004	M_{13}	0.0055
Single, 1 child (11-15)	0.0004	M_{14}	0.0021
Single, 2 children (Under 5)	0.0016	M_{15}	0.0027
Single, 2 children (5-10)	0.0009	M_{16}	0.0024
Single, 2 children (11-15)	0.0004	M_{17}	0.0087
Single, 3 children (Under 5)	0.0025	M_{18}	0.0004
Single, 3 children (5-10)	0.0009	M_{19}	0.00001
Single, 3 children (11-15)	0.0004	M_{20}	$\approx$0
Single, adult children	0.0001	M_{21}	0.0002
Single, no children	0.0025	M_{22}	$\approx$0
Part-time work multiplier, married	0.0003	M_{23}	0.00002
"	"	M_{24}	0.00004
Part-time work multiplier, single	0.0002	M_{25}	0.00003
"	"	M_{26}	0.00001
Initial log wage	0.00004	M_{30}	0.00003
s.d. of initial heterogeneity	0.00003	M_{39}	$\approx$0
s.d. of wage shocks	3×10^{-5}	M_{49}	3×10^{-6}
		M_{50}	1×10^{-6}
		M_{40}	3×10^{-6}
		M_{41}	1×10^{-6}
		M_{42}	$\approx$0
		M_{43}	1×10^{-6}
		M_{44}	$\approx$0
		M_{45}	1×10^{-6}
		M_{46}	1×10^{-6}
		M_{47}	2×10^{-5}
Persistence of wage process	8×10^{-5}	M_{48}	$\approx$0
		M_{40}	$\approx$0
		M_{41}	6×10^{-6}
		M_{42}	2×10^{-5}
		M_{43}	2×10^{-5}
		M_{44}	8×10^{-6}
		M_{45}	$\approx$0
		M_{46}	8×10^{-6}
		M_{47}	$\approx$0
Returns to full-time work	9×10^{-8}	M_{27}	1×10^{-8}

		M_{31}	8×10^{-8}
		M_{32}	7×10^{-8}
		M_{33}	1×10^{-8}
		M_{34}	2×10^{-9}
		M_{35}	8×10^{-9}
		M_{36}	≈ 0
		M_{37}	≈ 0
		M_{38}	≈ 0
Returns to part-time work	2×10^{-7}	M_{29}	≈ 0
Decline in experience from age	1×10^{-10}	M_{28}	≈ 0
		M_{31}	7×10^{-11}
		M_{32}	≈ 0
		M_{33}	≈ 0
		M_{34}	≈ 0
		M_{35}	≈ 0
		M_{36}	7×10^{-11}
		M_{37}	1×10^{-11}
		M_{38}	2×10^{-11}
Returns to not working	1×10^{-5}	M_{27}	2×10^{-6}
Probability of non-market childcare	0.0004	M_{51}	0.0084

Moment	Number	in Data	in Model
$Cov(w_{i,t}, w_{i,t-1})$	M_{48}	0.0680	0.0701
		(0.0660,0.0700)	
$s.d(\Delta w_{i,t})$	M_{49}	0.1993	0.1995
		(0.1954,0.2032)	
$Cov(\Delta w_{i,t}, \Delta w_{i,t-1})$	M_{50}	-0.0143	-0.0146
		(-0.0155,-0.0131)	
Proportion of Working Mothers using Childcare	M_{51}	0.277	0.278
		(0.259,0.295)	

Table A.8 Other Moments

Confident intervals listed under data estimates in brackets.

B Appendix to: Adult social care funding and its interaction with saving for retirement

B.1 Solution of the Model

B.1.1 Use of the Endogenous Grid Method

The estimation and simulation of the model requires the solution of household decisions across the model. Since there is no analytical solution to the problem, I solve the problem numerically using backwards recursion, starting from the latest possible age of death at 100.

For efficiency, I use the endogenous grid method to estimate saving and consumption decisions, following Carroll (2006). However, the model involves simultaneously solving for consumption and labour supply choices, with discrete labour supply and continuous consumption choices. The use of endogenous grid method to solve models with both discrete and continuous choices is much harder, due to discrete choices creating kinks in the value function. First, it requires calculating the Euler equation for every possible discrete labour supply choice in the first instance. This removes the risk of kinks appearing in the same period as the simultaneous choice, as optimal saving decisions are made for each possible labour supply choice.

However, in this setup, you might also expect kinks in the marginal utility of consumption in future periods. The reason they might occur is that a change in present-day consumption might change future labour supply decision, which would in turn change the marginal utility of consumption in the opposite direction. This means the marginal utility of wealth is not smooth, which means that attempting to solve the model using Euler equations creates problems. To get around this, I follow the solution proposed by Fella (2014). This involves finding the bounds of the assets space in which the consumption choice involves kinks in the marginal utility of wealth, and removing the non-optimal solutions within those bounds. In this way, I obtain a solution to the model.

B.1.2 Details of Solution and Simulation

I obtain the solution by using a discrete grid to represent the state space, across which I then interpolate upon simulating the population. I use a grid of 13 points for non-housing assets, 9 points for housing assets, and 4 points for the wage grid. In addition, the state space includes 5-point grids for men's and women's health, and a binary grid for whether households have access to family members who can help with care. Because they are mutually exclusive from other elements of the state space, I append one extra grid point on the wage grid, which represents those who have retired.

In each period, for each plausible point in the grid, I calculate the optimal consumption choice for each possible discrete choice. For those who are not already retired and own a home, there are six discrete choices: retiring or not retiring, each paired with each of keeping your home, selling your home, or downsizing. [1] This takes into account expectations for future periods over possible shocks to health and social care need. Having calculated optimal consumption decisions for each retirement and housing decision, I then choose the discrete choice which maximises the value function for each given consumption choice. I then input this into the value function, and iterate backwards onto the period representing a household one year earlier. I iterate this process starting from the primary resident being aged 100 down to the age of 50, until I have a solution grid of optimal consumption choices for each retirement and housing choice in the state space for each age, and the value of the value function for each of these discrete choices at each point.

I use this set of optimal choices to simulate the model. I apply the solution grid to the initial ELSA survey as described in Section 2.4 I then apply the solution grid to the simulations. I do this by interpolating across up to three dimensions (non-housing assets, housing, and wages), imputing consumption, housing, and retirement decisions from the shocks and state variables of the simulated individuals, by linear interpolation.

B.2 List of instruments for self-reported health

Condition	Coefficient	Standard Error
High blood pressure	0.24	(0.01)
Angina	0.33	(0.02)

[1] For those who are already retired or do not own their own home, there are fewer possible discrete choices available.

Myocardial infarction or coronary thrombosis	0.26	(0.02)
Congestive heart failure	0.37	(0.05)
Heart murmur	0.13	(0.02)
Abnormal heart rhythm	0.24	(0.01)
Diabetes or high blood sugar	0.46	(0.01)
Cerebral vascular disease	0.40	(0.02)
Chronic lung disease	0.66	(0.02)
Asthma	0.29	(0.01)
Arthritis	0.45	(0.01)
Osteoporosis	0.38	(0.01)
Cancer	0.33	(0.02)
Parkinson's	1.00	(0.05)
Any emotional, nervous, or psychiatric problems	0.41	(0.01)
Alzheimer's	0.35	(0.09)
Dementia	0.57	(0.04)
Glaucoma	0.13	(0.02)
Diabetic eye disease	0.20	(0.03)
Macular degeneration	0.09	(0.02)
Cataracts	0.12	(0.01)
Deafness	0.46	(0.02)

Table B.1 Link between different diseases and self-reported health

B.3 Model Fit

B.3.1 List of Moments

In this section I detail the fit of the model, and the specification of the moments used. I compare the moment values in the data and the model, and compare to the confidence intervals in the data. In each case the confidence intervals come from bootstrapping the moments in the data.

The moments show the that the model fits the data very well. All but three moments are within the confidence intervals of the data, and those that are not - the retirement rate at age 62, and home ownership rates for single and married people at age 77. Even for these, each is within three standard errors of the value in the data. The full list of moments is listed below, with confidence intervals for each moment in the data listed in brackets.

Moment	Number	in Data	in Model
Mean Log Annual Income, Age 57	M_1	10.23	10.24
		(10.19,10.27)	
Mean Log Annual Income, Age 62	M_2	9.98	9.99
		(9.94,10.02)	
Mean Log Annual Income, Age 67	M_3	9.45	9.43
		(9.37,9.53)	
Proportion Retired, Age 57	M_4	0.142	0.132
		(0.126,0.158)	
Proportion Retired, Age 62	M_5	0.315	0.341
		(0.297,0.333)	
Proportion Retired, Age 67	M_6	0.680	0.680
		(0.662,0.698)	
Home Ownership Rate, Single, Age 67	M_7	0.667	0.636
		(0.636,0.698)	
Home Ownership Rate, Single, Age 77	M_8	0.637	0.686
		(0.604,0.670)	
Home Ownership Rate, Single, Age 87	M_9	0.652	0.643
		(0.607,0.697)	
Home Ownership Rate, Married, Age 67	M_{10}	0.899	0.902
		(0.885,0.913)	
Home Ownership Rate, Married, Age 77	M_{11}	0.879	0.900
		(0.863,0.895)	
Home Ownership Rate, Married, Age 87	M_{12}	0.895	0.861
		(0.860,0.930)	
Mean Log Total Assets, Age 85	M_{13}	5.052	5.106
		(4.944,5.160)	
Mean Log Care Spend	M_{14}	8.061	
		(7.999,8.233)	

Table B.2 Moments in Data and Model

B.3.2 Relevance of Moments

I now outline the importance of different moments for determining value of each parameter. I do this by using methods pioneered by Andrews et al. (2017) and Honore et al. (2020). In particular, I measure how sensitive the standard errors of the parameter values are to removing each moment from their estimation. By measuring the change in the standard errors when a moment is removed, I gain a sense of how significant a moment is to that parameter: if the standard error of the parameter would double were the moment not there, then that moment can be considered as very important for the estimation of that parameter. Due to space, I do not report all combinations of moments and parameters, but only those that might be intuitively significant to each parameter.

In the case of almost every parameter, removing the intuitively significant moments would at least double the variance of the estimate. This indicates that the estimation is well specified, as the parameters are predominantly determined by the moments one would expect them to. The one exception to this is the preference for bequests, where the inclusion of total assets at age 85 as a moment only slightly reduces the variance, indicating that the preference for bequests is less well identified than other parameters.

Parameter	Variance	Moment Removed	Change in Variance
Linear Log Wage Change With Age	2×10^{-6}	M_1	4×10^{-7}
		M_2	4×10^{-7}
		M_3	1×10^{-6}
Quadratic Log Wage Change With Age	4×10^{-8}	M_1	≈ 0
		M_2	9×10^{-9}
		M_3	3×10^{-8}
Utility Cost of Work, Age 50	5×10^{-8}	M_4	1×10^{-7}
Change in Cost of Work Each Year	4×10^{-10}	M_5	7×10^{-11}
		M_6	≈ 0
Utility from Home Ownership, Single	3×10^{-9}	M_7	3×10^{-9}
		M_8	4×10^{-10}
		M_9	1×10^{-9}
Utility from Home Ownership, Married	9×10^{-10}	M_{10}	≈ 0
		M_{11}	3×10^{-11}
		M_{12}	6×10^{-8}
Preference for Bequests	0.0196	M_{13}	0.0008

Preference for Care	0.0025	M_{14}	0.0292

C Appendix to: Adult social care funding and its interaction with saving for retirement

C.1 Standard difference-in-differences results

Table C.1 Impact of the EMA on education participation aged 16-17 (FSM students)

	Baseline	Standard DiD
Full Time Education	0.629	0.036**
		(0.016)
Further Education	0.539	0.018
		(0.021)
In School	0.114	0.017
		(0.012)
A-Levels	0.204	0.002
		(0.010)
In Training	0.081	-0.013**
		(0.005)
NEET	0.281	-0.022
		(0.015)
Observations		89,220
LAs $\times$ Years		1440

Age 16-17 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

Table C.2 Impact of the EMA on education participation aged 17-18 (FSM students)

	Baseline	Standard DiD
Full Time Education	0.500	0.025
		(0.015)
Further Education	0.444	0.022
		(0.020)
In School	0.066	0.005
		(0.006)
A-Levels	0.182	0.004
		(0.010)
In Training	0.104	-0.011
		(0.009)
NEET	0.338	-0.002
		(0.015)
Observations		89,220
LAs $\times$ Years		1440

Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

Table C.3 Impact of the EMA on education participation by gender (FSM students)

	Standard DiD (Age 16-17)	Standard DiD (Age 17-18)
Boys' Participation	0.057**	0.023
	(0.021)	(0.020)
Observations	44,943	44,943
Girls' Participation	0.015	0.027
	(0.019)	(0.020)
Observations	44,275	44,275
LAs $\times$ Years	180	180

Age 16-17 and Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

Table C.4 Impact of the EMA on education participation by prior academic performance (FSM students)

	Standard DiD (Age 16-17)	Standard DiD (Age 17-18)
5 A*-C Participation	-0.016	-0.022
	(0.020)	(0.020)
Observations	21,702	21,702
Not 5 A*-C Participation	0.052**	0.039*
	(0.022)	(0.021)
Observations	67,516	67,516
LAs × Years	180	180

Age 16-17 and Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

Table C.5 Impact of the EMA on education participation by special educational needs (FSM students)

	Standard DiD (Age 16-17)	Standard DiD (Age 17-18)
SEN Participation	0.035	0.043
	(0.028)	(0.026)
Observations	31,117	31,117
No SEN Participation	0.037**	0.019*
	(0.017)	(0.016)
Observations	58,108	58,108
LAs × Years	180	180

Age 16-17 and Age 17-18 represent the time in the academic year covered by these two ages. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

Table C.6 Impact of the EMA on work at age 17 (FSM students)

	Baseline	Standard DiD
Annual Earnings, Age 17, £	1634	-173*
		(100)
% Earning above LEL, Age 17	0.118	-0.016
		(0.011)
Earnings above 0, Age 17	0.381	-0.020
		(0.017)
Observations		89,220
Annual Earnings, Age 17, £	1122	-116**
In Education		(43)
Observations		65,390
LAs × Years		1440

Age 17 represents the tax year (April-April) in which everyone in the cohort is aged 17 at some point. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

Table C.7 Impact of the EMA on qualifications (FSM students)

	Baseline	Standard DiD
A Level Results, Standardised	-1.010	0.096
		(0.071)
Observations		25,150
Level 1 Qualification Or Above	0.711	0.003
		(0.017)
Level 2 Qualification Or Above	0.366	0.011
		(0.016)
Level 3 Qualification Or Above	0.161	0.008
		(0.010)
University Attendance	0.084	0.007
		(0.008)
Observations		89,220
LAs $\times$ Years		1440

Level qualifications considered by age 19, and only include those obtained at FE colleges, not schools. University degrees considered by age 23. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

Table C.8 Impact of the EMA on earnings (FSM students)

	Baseline	Standard DiD
	Age 22	
Earnings (Log)	£8200	0.011
		(0.025)
	Age 24	
Earnings (Log)	£9600	-0.028
		(0.019)
	Age 26	
Earnings (Log)	£10,500	-0.038
		(0.027)
	Age 28	
Earnings (Log)	£11,400	-0.002
		(0.022)

Results for those who have annual earnings greater than zero in the given year. *,** and *** indicate significance at the 10%, 5%, and 1% level respectively.

References

Adams, Abigail, Mathias Jensen, and Barbara Petrongolo, "Birth Timing and Spacing: Implications for Parental Leave Dynamics and Child Penalties," Economics Series Working Papers 1048, University of Oxford, Department of Economics July 2024.

Adda, J, C Dustmann, and K Stevens, "The Career Costs of Children," *Journal of Political Economy*, 2017, *125* (2), 293–337.

Alzúa, M.L., Guillermo Cruces, and Carolina Lopez, "Long-Run Effects Of Youth Training Programs: Experimental Evidence From Argentina," *Economic Inquiry*, October 2016, *54* (4), 1839–1859.

Andrews, I, M Gentzkow, and J. M. Shapiro, "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *The Quarterly Journal of Economics*, 2017, *132* (4), 1553–1592.

Arkhangelsky, Dmitry, Susan Athey, David A. Hirshberg, Guido W. Imbens, and Stefan Wager, "Synthetic Difference-in-Differences," *American Economic Review*, December 2021, *111* (12), 4088–4118.

Attanasio, O and G Weber, "Is Consumption Growth Consistent With Intertemporal Optimization?," *Journal of Political Economy*, 1995, *103* (6), 1121–1157.

_ , Hamish Low, and V Sánchez-Marcos, "Explaining changes in female labor supply in a life-cycle model," *American Economic Review*, September 2008, *98* (4), 1517–1552.

Attanasio, Orazio, Costas Meghir, and Ana Santiago, "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA," *Review of Economic Studies*, 2012, *79* (1), 37–66.

Baird, Sarah, Craig McIntosh, and Berk Özler, "Cash or Condition? Evidence from a Cash Transfer Experiment," *The Quarterly Journal of Economics*, 2011, *126* (4), 1709–1753.

_ , Joan Hamory Hicks, Michael Kremer, and Edward Miguel, "Worms at Work: Long-run Impacts of a Child Health Investment," *The Quarterly Journal of Economics*, 2016, *131* (4), 1637–1680.

Barczyk, Daniel and Matthias Kredler, "Evaluating Long-Term-Care Policy Options, Taking the Family Seriously," *The Review of Economic Studies*, 2018, *85* (2), 766–809.

_ **, Sean Fahle, and Matthias Kredler**, "Save, Spend, or Give? A Model of Housing, Family Insurance, and Savings in Old Age," *The Review of Economic Studies*, 2023, *90* (5), 2116–2187.

Barham, Tania, Karen Macours, and John A. Maluccio, "More Schooling and More Learning?: Effects of a Three-Year Conditional Cash Transfer Program in Nicaragua after 10 Years," IDB Publications (Working Papers) 4584, Inter-American Development Bank July 2013.

Barrera-Osorio, Felipe and Deon Filmer, "Incentivizing Schooling for Learning: Evidence on the Impact of Alternative Targeting Approaches," *Journal of Human Resources*, 2016, *51* (2), 461–499.

_ **, Leigh L. Linden, and Juan E. Saavedra**, "Medium- and Long-Term Educational Consequences of Alternative Conditional Cash Transfer Designs: Experimental Evidence from Colombia," *American Economic Journal: Applied Economics*, July 2019, *11* (3), 54–91.

_ **, Marianne Bertrand, Leigh L. Linden, and Francisco Perez-Calle**, "Improving the Design of Conditional Transfer Programs: Evidence from a Randomized Education Experiment in Colombia," *American Economic Journal: Applied Economics*, April 2011, *3* (2), 167–195.

Behrman, Jere R., Susan W. Parker, and Petra E. Todd, "Do Conditional Cash Transfers for Schooling Generate Lasting Benefits?: A Five-Year Followup of PROGRESA/Oportunidades," *Journal of Human Resources*, 2011, *46* (1), 93–122.

Bensnes, Simon, Ingrid Huitfeldt, and Edwin Leuven, "Reconciling Estimates of the Long-Term Earnings Effect of Fertility," IZA Discussion Papers 16174, Institute of Labor Economics (IZA) May 2023.

Bernal, Raquel, "The Effect Of Maternal Employment And Child Care On Children'S Cognitive Development," *International Economic Review*, November 2008, *49* (4), 1173–1209.

Blundell, R, J Britton, M Costa Dias, and E French, "The impact of health on labor supply near retirement," IFS Working Papers W17/18, Institute for Fiscal Studies 2017.

_ **, M Browning, and C Meghir**, "Consumer Demand and the Life-Cycle Allocation of Household Expenditures," *The Review of Economic Studies*, 1994, *61* (1), 57–80.

_ **, M Costa Dias, Costas Meghir, and J Shaw**, "Female labour supply, human capital and welfare reform," *Econometrica*, 2016, *84* (5), 1705–1753.

Boca, Daniela Del, Christopher Flinn, and Matthew Wiswall, "Household Choices and Child Development," *The Review of Economic Studies*, 2014, *81* (1), 137–185.

Bolton, Paul, "Education Maintenance Allowance (EMA) Statistics," Technical Report SNSG/5778, House of Commons Library 2011.

Brewer, Mike, Sarah Cattan, Claire Crawford, and Birgitta Rabe, "Does more free childcare help parents work more?," *Labour Economics*, 2022, *74* (C).

Bueren, Jesus, "Long-Term Care Needs and Savings in Retirement," *Review of Economic Dynamics*, July 2023, *49*, 201–224.

Bögl, Sarah, Jasmin Moshfegh, Petra Persson, and Maria Polyakova, "The Economics of Infertility: Evidence from Reproductive Medicine," NBER Working Papers 32445, National Bureau of Economic Research, Inc May 2024.

Carroll, C.D., "The Method of Endogenous Gridpoints for Solving Dynamic Stochastic Optimization Problems," *Economics Letters*, 2006, *91* (3), 312–320.

Carta, Francesca and Lucia Rizzica, "Early kindergarten, maternal labor supply and children's outcomes: Evidence from Italy," *Journal of Public Economics*, 2018, *158*, 79–102.

Casarico, Alessandra and Salvatore Lattanzio, "Behind the child penalty: understanding what contributes to the labour market costs of motherhood," *Journal of Population Economics*, July 2023, *36* (3), 1489–1511.

Chan, Marc K. and Kai Liu, "Life-cycle and intergenerational effects of child care reforms," *Quantitative Economics*, July 2018, *9* (2), 659–706.

Crawford, R and C O'Dea, "Household portfolios and financial preparedness for retirement," *Quantitative Economics*, 2020, *11*, 637–670.

Daycare Trust, "Childcare Costs Survey 2011," February 2011.

Dearden, Lorraine and Alexandra Heath, "Income support and staying in school: what can we learn from Australia's AUSTUDY experiment?," *Fiscal Studies*, 1996, *17* (4), 1–30.

_ , Carl Emmerson, Christine Frayne, and Costas Meghir, "Conditional cash transfers and school dropout rates," *Journal of Human Resources*, 2009, *44* (4), 827–857.

Dobrescu, Loretti I., "To Love or to Pay: Savings and Health Care in Older Age," *Journal of Human Resources*, 2015, *50* (1), 254–299.

Eckstein, Z, M Keane, and O Lifshitz, "Career and Family Decisions: Cohorts born 1935-1975," *Econometrica*, 2019, *87* (1), 217–253.

Engen, E. M., W. G. Gale, and C. E. Uccello, "The Adequacy of Household Saving," *Brookings Papers on Economic Activity*, 1999, pp. 65–187.

Evans, David K and Mũthoni Ngatia, "School Uniforms, Short-Run Participation, and Long-Run Outcomes: Evidence from Kenya," *The World Bank Economic Review*, 2021, *35* (3), 705–719.

Fadlon, Itzik and Torben Heien Nielsen, "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records," *American Economic Journal: Applied Economics*, July 2021, *13* (3), 1–30.

Fella, G, "A Generalized Endogenous Grid Method for Non-Smooth and Non-Concave problems," *Review of Economic Dynamics*, 2014, *17* (2), 329–344.

Filippucci, Francesco, "What Do NEETs Need? The Effect of Combining Activation Policies and Cash Transfers," 2022.

Filmer, Deon and Norbert Schady, "Does more cash in conditional cash transfer programs always lead to larger impacts on school attendance?," *Journal of Development Economics*, September 2011, *96* (1), 150–157.

French, E and J.B. Jones, "The effects of health insurance and self-insurance on retirement behavior," *Econometrica*, 2011, *79* (3), 693–732.

Galiani, Sebastian and Patrick J. McEwan, "The heterogeneous impact of conditional cash transfers," *Journal of Public Economics*, 2013, *103* (C), 85–96.

García, Sandra and Juan Saavedra, "Educational Impacts and Cost-Effectiveness of Conditional Cash Transfer Programs in Developing Countries: A Meta-Analysis," NBER Working Papers 23594, National Bureau of Economic Research, Inc July 2017.

Glewwe, Paul and Ana Lucia Kassouf, "The impact of the Bolsa Escola/Familia conditional cash transfer program on enrollment, dropout rates and grade promotion in Brazil," *Journal of Development Economics*, 2012, *97* (2), 505–517.

Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti, "The parental pay gap over the life cycle: Children, jobs, and labor supply," *Journal of Economic Dynamics and Control*, 2024, *169* (C).

Heckman, James, Seong Hyeok Moon, Rodrigo Pinto, Peter Savelyev, and Adam Yavitz, "Analyzing social experiments as implemented: A reexamination of the evidence from the HighScope Perry Preschool Program," *Quantitative Economics*, July 2010, *1* (1), 1–46.

Honore, B, T Jørgensen, and A de Paula, "The Informativeness of Estimation Moments," *Journal of Applied Econometrics*, 2020, *35* (7), 797–813.

Hubbard, R.G., J Skinner, and S.P Zeldes, "Precautionary Saving and Social Insurance," *Journal of Political Economy*, 1995, *103* (2), 360–399.

Jones, L.E., R.E. Manuelli, and E.R. McGrattan, "Why Are Married Women Working So Much?," *Journal of Demographic Economics*, 2015, *81* (1), 75–114.

Kaplan, G and G. L. Violante, "A Model of the Consumption Response to Fiscal Stimulus Payments," *Econometrica*, 2014, *82* (4), 1199–1239.

Kim, Kimin and Myoung jae Lee, "Difference in differences in reverse," *Empirical Economics*, 2019, *57*, 705–725.

Kleven, H, C Landais, and J.E. Søgaard, "Children and Gender Inequality: Evidence from Denmark," *American Economic Journal: Applied Economics*, 2019, *11* (4), 181–209.

Kleven, Henrik, Camille Landais, and Gabriel Leite-Mariante, "The Child Penalty Atlas," NBER Working Papers 31649, National Bureau of Economic Research, Inc August 2023.

Kopecky, Karen A. and Tatyana Koreshkova, "The Impact of Medical and Nursing Home Expenses on Savings," *American Economic Journal: Macroeconomics*, July 2014, *6* (3), 29–72.

Kreider, B and J. V. Pepper, "Disability and Employment: Reevaluating the Evidence in Light of Reporting Errors," *Journal of the American Statistical Association*, 2007, *102* (478), 432–441.

Kuka, Elira and Na'ama Shenhav, "Long-Run Effects of Incentivizing Work after Childbirth," *American Economic Review*, June 2024, *114* (6), 1692–1722.

Low, Hamish and Luigi Pistaferri, "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off," *American Economic Review*, 2015, *105* (10), 2986–3029.

_ , Costas Meghir, Luigi Pistaferri, and Alessandra Voena, "Marriage, Labour Supply and the Dynamics of the Social Safety Net," Working Paper W18/05, IFS 2018.

Molina-Millan, Teresa, Tania Barham, Karen Macours, John A. Maluccio, and Marco Stampini, "Long-term Impacts of Conditional Cash Transfers in Latin America: Review of the Evidence," IDB Publications (Working Papers) 7891, Inter-American Development Bank October 2016.

Nardi, M De, E French, J.B. Jones, and R McGee, "Why Do Couples and Singles Save During Retirement," NBER Working Papers 28828, National Bureau of Economic Research 2021.

Nardi, Mariacristina De, Eric French, and John Bailey Jones, "Medicaid Insurance in Old Age," *American Economic Review*, November 2016, *106* (11), 3480–3520.

Nollenberger, N and N Rodríguez-Planas, "Full-time universal childcare in a context of low maternal employment: quasi-experimental evidence from Spain," *Labour Economics*, 2015, *36*, 124–136.

Office for National Statistics, "Care homes and estimating the self-funding population, England: 2019 to 2020."

Riccio, James, Nadine Dechausay, Cynthia Miller, Stephen Nuñez, Nandita Verma, and Edith Yang, *Conditional Cash Transfers in New York City*, New York: MDRC, 2013.

Sabates, Ricardo and Leon Feinstein, "Effects of government initiatives on youth crime," *Oxford Economic Papers*, 2008, *60* (3), 462–483.

Schochet, Peter Z., "Statistical Power for Random Assignment Evaluations of Education Programs," *Journal of Educational and Behavioral Statistics*, March 2008, *33* (1), 62–87.

Scholz, J. K., A. Seshadri, and S. Khitatrakun, "Are Americans saving "optimally" for retirement?," *Journal of Political Economy*, 2006, *114*, 607–643.

Sefton, J, J van de Ven, and M Weale, "Means Testing Retirement Benefits: Fostering Equity or Discouraging Savings?," *The Economic Journal*, 2008, *118* (528), 556–590.

Stern, S, "Measureing the Effect of Disability on Labor Force Participation," *The Journal of Human Resources*, 1989, *24* (3), 361–395.

Tartari, Melissa, "Divorce And The Cognitive Achievement Of Children," *International Economic Review*, May 2015, *56* (2), 597–645.