Household Decisions Over the Life Cycle

Timothy Hunt

Nuffield College

University of Oxford

April 2026

A thesis submitted for the degree of

Doctor of Philosophy in Economics

Acknowledgements

I am deeply grateful to my academic supervisors Hamish Low, Abi Adams and Romuald Méango. In particular, I thank Hamish and Abi for their guidance since the early stages of the DPhil and their willingness to continue with supervision even when circumstances took them elsewhere.

Many other people in the Oxford Economics Department have been important sources of advice, particularly Miguel Ballester, Péter Eső, James Forder and Ian Jewitt.

I would not have been able to complete the thesis without the generous support of Nuffield College and the guidance of my college supervisor Ian Crawford.

My academic stay at the University of Michigan in 2024 was an invaluable experience and my work benefitted greatly from discussions with my host, Ana Reynoso, as well as others including Emilio Borghesan, Thomas Crossley and Mel Stephens.

Friends at Oxford and elsewhere, and family back at home, have been constant sources of support throughout the long process of completing the DPhil.

Finally, I thank my partner Ali for her patience and encouragement, and dedicate this thesis to her and the good old days to come.

Table of Contents

Abstract

This thesis features three chapters on distinct questions in the economics of ageing.

Chapter 1 examines the provision of informal care to elderly parents in the US. Using a dynamic model of sibling interaction, it shows that observed differences in opportunity costs explain little of the gender gap in caregiving, while there is an important role for strategic interaction in exacerbating the underlying gender care gap. These findings imply that policies aimed at closing the gender wage gap may not have significant effects on the gender care gap.

Chapter 2 analyses how households respond to mistakes in retirement planning. It exploits data on variation in beliefs about the State Pension Age in England and builds a a life-cycle model with rational inattention over pension policies. The results show that the welfare costs of inattention in this setting are heterogeneous but modest.

Chapter 3 studies the interaction between long-term care policy and the housing market. The focus is on the "homestead exemption" in means-tested support with long-term care costs, under which housing wealth is better protected from these costs. An overlapping-generations model of the housing market shows that this exemption distorts housing demand and raises house prices. A budget-balanced repeal of this policy brings significant welfare benefits.

0 | Introduction

This thesis presents three chapters on different research questions in the economics of ageing and household decisions over the life cycle. The overarching aim of the thesis is to analyse the degree to which the family and the state are able to mediate the shocks that people face late in their lives and the effects of these efforts on the rest of the population.

Chapter 1, *Sharing the caring? Dynamic interaction between siblings in the provision of care to parents*, considers the issue of informal care by adult children to their elderly parents. In the US setting of the paper, daughters provide around 3 times the monthly care hours of sons to parents with care needs. In principle, this gender care gap could be driven by observable differences in opportunity costs of providing care, such as wage differentials in the labour market, or it could be driven by unobservable preference differences between sons and daughters, whether those differences be understood as preferences on the part of the children, of the parents, or of society as a whole. The chapter seeks to address two research questions: how much of the gender care gap is explained by observed differences in opportunity costs versus unobserved differences in preferences? And what is the role of strategic interaction between siblings in exacerbating this gender care gap?

The paper first uses Health and Retirement Study (HRS) data to present some descriptive evidence on the topic of informal caregiving by adult children in the US. Important results include evidence of strategic interaction between siblings, in the sense that siblings change their care provision depending on the characteristics of their siblings, and a lack of evidence of substantial compensation given to children for their caregiving, motivating a non-cooperative model of sibling interaction.

I then develop a dynamic discrete-choice model of interaction between siblings in the provision of care to parents. The model is a game where the two players, a younger and an older child, make decisions over labour force participation, location and care provision every period. Care to parents is modelled as a public good: both children enjoy the benefits of having a cared-for parent but find this costly to provide. There is rich heterogeneity by demographics in the preferences for and costs of providing care.

I estimate the model by maximising the pseudo-likelihood of observed choices and use the estimated model to run some counterfactual experiments. I find that observed differences in wages explain very little of the gender care gap. This is consistent with reduced-form evidence that within families (i.e. using family fixed effects) it is the most educated children - thus likely those with the greatest opportunity costs of providing care - who are the most likely to provide care, casting doubt on the opportunity-cost explanation of the gender care gap. I also find that strategic interaction between siblings explains around 31% of the gender care gap. These results suggest that i) policies directed at closing the gender wage gap will likely not have a significant impact on the gender care gap and ii) policies which do successfully change individual agents' incentives to provide care will, through strategic interaction, also have important effects on the care provision by other family members.

Chapter 2, *Where's my pension gone? Labour supply effects of mistakes in retirement planning*, examines how households respond to sub-optimal planning for retirement. In general, this is a difficult issue to analyse because households who, for instance, do not save much for retirement might be planning optimally but just have high discount rates or resources unobserved by the econometrician. I make progress by using English data on people's beliefs about their State Pension Age (SPA), the age at which they are eligible to receive their public pension, which makes up a significant proportion of retired households' income. By observing how agents change their behaviour after switching from having incorrect beliefs to correct beliefs about their SPA I am able to isolate how people respond to the realisation of genuine mistakes in retirement planning.

In particular, I exploit an increase in the SPA for women, announced in 1995 and enacted from 2010. Using data from the English Longitudinal Study of Ageing (ELSA) I first document the correlates of having incorrect beliefs about one's SPA and I present some descriptive event studies showing how switching to having correct beliefs about one's SPA is associated with an increase in household labour supply but no notable change in other decision variables. I argue, however, that it is implausible to treat the event of correcting one's beliefs about the SPA as being an exogenous shock insofar as there is evidence that many women were aware of the reform in general and just were unable to answer correctly how it applied to them. I use these facts to motivate a life-cycle model of over-50s labour supply and consumption with a rational inattention component: agents rationally decide whether or not to pay the cost of becoming informed about their true SPA, and where this cost of paying attention (as well as costs pertaining to their labour force decision) depends on their characteristics.

I estimate the model via the Method of Simulated Moments, and use the estimated model to carry out welfare analysis and counterfactual experiments. Agents in the model increase their labour supply once they realise they have made mistakes in retirement planning. However, I find that the cost of starting the model inattentive is small: on average ex post, households who start the model with incorrect beliefs would be indifferent between only a £457 increase in wealth at the start of the model and living in a world where they had always had correct beliefs, relative to annual median household consumption at age 60 in the model of £20.3k. There is, however, notable heterogeneity in the costs of inattention by household characteristics, with those households with less education or closer to retirement at the start of the model facing higher costs. I also analyse potential policies to mitigate the cost of inattention, finding that neither wage subsidies for re-entering work nor increased welfare payments are particularly effective in this regard.

Chapter 3, *Long-term care policy and housing market efficiency*, considers the effects of government policies to protect the wealth of older people with care needs. In many countries, including this chapter's UK setting, the government provides means-tested support with long-term care needs where housing wealth is treated differently to financial wealth in the means test. In particular, housing wealth is generally not counted as part of the means test (the "homestead exemption"). This means that housing is more protected against long-term care costs than financial wealth is, distorting older people's housing consumption upwards and creating an incentive against downsizing late in life. This chapter analyses the degree to which such "homestead exemption" policies affect the housing market and the distribution of welfare in society.

After first presenting some policy context regarding housing consumption over the life cycle, I set out a dynamic model of the housing market. The model features overlapping generations of households, each of which every period makes consumption and housing decisions. Households face costs of moving house and experience persistent utility shocks for particular houses. Households also suffer health shocks resulting in long-term care costs and receive means-tested government support funded through income tax and a transaction tax on housing. The market clearing housing price is determined endogenously by demand and (fixed) supply, where housing demand is driven by preferences for housing as a consumption good and as a relatively protected store of wealth.

The model is estimated by the Method of Simulated Moments. I show using the estimated model that a budget-balanced repeal of the homestead exemption would reduce house prices by 6.5%, increase the housing consumption of the young, and boost average welfare by the equivalent of a £144 (0.6%) increase in consumption per year. Those who benefit the most are people without much housing wealth in the initial steady state and those who benefit the least are those with health problems or who have inherited more money in the initial steady state. The chapter thus points to important potential welfare gains from relaxing the homestead exemption, even though these welfare gains will be unevenly distributed and some households will lose out.

The thesis ends with a Conclusion that draws together the main findings from the three chapters and points towards directions for future research.

1 | Sharing the caring? Dynamic interaction between siblings in the provision of care to parents

Sharing the caring? Dynamic interaction between siblings in the provision of care to parents

Abstract

Adult daughters provide nearly three times as many hours of care to elderly parents as adult sons do. I analyse the role of strategic interaction between siblings in exacerbating this gap in care provision. To do so, I build and estimate a dynamic discrete-choice game in which siblings make care, location and work choices. I find that the opportunity for strategic play increases gender differences in caring responsibilities. Sons in particular strategically shirk providing care as they believe their sibling is relatively likely to provide care in their absence. Counterfactual experiments show that if siblings instead took care, location and work choices independently then the gender care gap would be around 31% smaller. Also, I find that observed differences in wages explain very little of the gender care gap, suggesting that policies to close the gender wage gap may make little difference to the gender care gap.

1.1 Introduction

Children are a key source of care for their elderly parents: in the US, around 17% of all adult children will provide care to a parent at some point in their lives and conditional on providing positive care hours an adult child will provide on average 77 hours of care per month (Wettstein and Zulkarnain 2017). Studies have shown the adverse effects of caregiving on labour force attachment (Van Houtven, Coe, and Skira 2013) and mental and physical health (Bom et al. 2019), and there is little evidence of caregivers receiving direct contemporaneous cash transfers in exchange for their care (Barczyk and Kredler 2018). It is generally daughters and not sons who select into caregiving (Grigoryeva 2017), with my calculations in this paper showing daughters providing around 3 times the care hours of sons even though all children plausibly benefit from the public good of having their parent cared for. Why do we observe such a gender gap in selecting into this role?

The existing literature has tried to analyse the issue of children providing care to their elderly parent(s) either from the perspective of an individual adult child (Ko 2022; Mommaerts 2024; Skira 2015) or with multiple adult children in a static setting (Byrne et al. 2009; Fontaine, Gramain, and Wittwer 2009), imposing different preferences or costs of caring for different types of children to match gaps in care provision. While this literature has yielded many rich insights, it has not accounted for the fact that care provision by children is fundamentally a problem of dynamic interaction. Modelling dynamics is crucial because children's decisions in response to parental health shocks will affect their ability to provide care in the future; modelling interactions is crucial because children make these decisions with the knowledge that their siblings are making similar calculations and because each child's decision affects the optimal behaviour of all their siblings. This creates a space for strategic play over time between siblings to amplify gaps arising from any difference in preferences or costs of caring.

A simple example will illustrate. Suppose Adam and Beth are siblings. Suppose there is something about Beth's preferences or costs for providing care that make her more likely to provide care (she has a lower opportunity cost of providing care, she finds it less burdensome to provide care, their common parent prefers Beth's care, and so on). If Adam and Beth made decisions independently, without regard for what the other was doing, then Beth would on average provide more care because of these differences in preferences and costs. However, if Adam knows Beth is likely to provide care, then other things being equal he will reduce his care effort; Beth, knowing this, will increase her care effort to compensate, meaning Adam could cut his care effort more, and so on. Thus, the possibility for strategic interaction can amplify existing gender care gaps. This amplification could be particularly important in dynamic games where agents' costs and preferences of providing care are not fixed exogenously but are instead endogenously generated by past decisions: Adam will be even less likely to provide care today if he sees today's care decision as locking him in to the caring role in future.

In this paper I explore this type of interaction by situating the decisions of adult children over the care provided to their elderly parents in a rich dynamic model, recognising a) that children may have different preferences over their parent's care arrangements and b) that children can take strategic actions to make it easier for them to provide (or avoid the burden of providing) care in future. In doing so, I make two key contributions to the literature on informal care provision within families. Firstly, I present the first estimates of the importance of strategic interaction in driving the gender care gap: specifically, a baseline estimate that the gender care gap would be 31% smaller for the population studied if there were no strategic interaction between siblings. Secondly, I provide new insight into the drivers of the gender care gap: differences in wages between sons and daughters explain very little of the gender care gap.

These findings have important policy implications. Given that the gender wage gap explains so little of the gender care gap, it is unlikely that policies to reduce the gender wage gap in isolation will make much difference to the gender care gap. However, given the importance of strategic interaction, policies which do manage to substantively change agents' incentives to provide care could significantly change the level and distribution of overall care provision as other agents strategically alter their caregiving in response.

The paper proceeds as follows. To motivate my structural approach, I first use the US Health and Retirement Study to document some stylised facts about the provision of care for elderly parents by their adult children. In keeping with previous results in the literature, I confirm that it is generally daughters who end up providing the care and I exploit variation in gender composition of sibling groups to offer some prima facie evidence that children's caregiving is not independent of the caregiving of their siblings, with care provided by sons being crowded out by care provided by daughters. However, in contrast to previous authors (Barczyk and Kredler 2018; Groneck 2017), I find little evidence of substantive compensation for most caregivers, either in the form of bequests or non-cash inter vivos transfers, which raises the question of what drives these people into the caregiving role.

I then use these stylised facts to motivate a dynamic model of non-cooperative strategic interaction between two adult siblings. Each period, the siblings simultaneously make location, labour and (if their common parent is sick) care decisions. Sons and daughters differ in their preferences for providing care and in their income from working versus not working. Children find it costly to provide care themselves but each child receives a benefit if any child provides care, meaning that the model resembles a public good provision problem. The children split the parent's bequest equally between them when the parent dies but otherwise face no financial incentive to provide care and cannot transfer money between themselves.

In keeping with the finding in Hiedemann, Sovinsky, and Stern (2018) about the importance of state dependence in caring arrangements, the model incorporates direct or implicit transition costs in changing one's location, labour force participation or caregiving, which in turn imply transition costs associated with changing caring arrangements. In particular, the children suffer a wage penalty if they have been absent from the labour market and they suffer a utility penalty from changing location or starting providing care. These transition costs are the key feature distinguishing this dynamic model from static models of care provision because care, location and work decisions today have a bearing on the optimal decisions in the next period.

I estimate the model by maximising the pseudo log-likelihood of the observed decisions (Aguirregabiria and Mira 2007). In particular, to make the estimation tractable, I use forward simulation to construct agents' value functions using the process set out in Bajari, Benkard, and Levin (2007) and employed in Ko (2022). Key estimated parameters include preference differences for providing care between sons and daughters and parameters capturing the extent to which adult children benefit from their siblings' care provision. I show that the model is able to match key patterns in the data and I carry out statistical tests that reject more parsimonious versions of the model.

I use the estimated model to conduct counterfactual experiments to quantify the drivers of the gender care gap: I examine the relative contributions of wages and preferences to this gap, and I assess the role of strategic interaction in exacerbating this gap. These exercises shed light on what we might expect to happen to the gender care gap in future given broader demographic trends such as a shrinking gender wage gap and rise in one-child families. I find that preferences are far more important than wages in driving the gender care gap, with imposing identical preferences between sons and daughters shrinking the gap by 81%, and I find that if every child assumed that their siblings would never provide care (so that there are no considerations of strategic interaction) then the gap would be 31% smaller.

This paper contributes to the literature on informal care arrangements within families (Byrne et al. 2009; Checkovich and Stern 2002; Fontaine, Gramain, and Wittwer 2009; Hiedemann, Sovinsky, and Stern 2018; Mommaerts 2024; Stern 2023), and in particular to the literature on dynamic strategic interaction over care provision within a family (Barczyk and Kredler 2018; Ko 2022; Sovinsky and Stern 2016). Previous papers on dynamic strategic interaction have focussed on interaction between a parent and child, rather than on strategic interaction between children; instead, to the best of my knowledge this has only been studied in a static setting (e.g. Byrne et al. (2009)). By modelling interaction between children in a rich dynamic setting I am able to capture the full importance of strategic play in driving the unequal distribution of the care burden.

Another important methodological contribution of this paper is the modelling of location choice along with caring and labour choices in a dynamic model of care provision.

Previous papers have treated child location as fixed (e.g. Ko (2022)) or have modelled the location decision of adult children as a one-shot decision taken early on in life (e.g. Stern (2023)). In practice, adult children can change their location in response to parental health shocks by moving back close to their parents in order to provide care and there is some evidence that they do so (e.g. Hiedemann, Sovinsky, and Stern (2018)). It is important to allow for this channel of behaviour to properly model a child's ability to alter their own cost of providing care and thus their ability to strategically shirk caregiving[1].

This paper is closest in approach to Ko (2022), which models dynamic interaction between a parent and child and uses this to explain low long-term care insurance takeup. I use many of the same modelling assumptions. For instance, I use a similar dynamic framework of non-cooperative interaction over caregiving decisions, combined with a labour participation decision on the part of the child. As suggested above, the key differences are that I model interaction between children, rather than between a parent and a child, in order to work out what drives differences in provision between children, and that I allow children's location to change endogenously. On the other hand, to make the model tractable, I abstract from modelling LTC insurance and in particular the interaction between the LTCI demand and the availability of informal care, a key focus of Ko (2022).

The inclusion of the labour force decision for children, and the emphasis on the costs associated with leaving the labour force to provide care, links this paper to Skira (2015), which presents a dynamic model of women's labour force participation and caregiving for an elderly parent. This paper expands on that paper by modelling the interaction between multiple children and allowing children choice along more dimensions, though this extra complexity comes at the cost of a simpler labour market and wage model than in Skira (2015).

The paper also draws from Hiedemann, Sovinsky, and Stern (2018), which similarly provides dynamic models of care provision within a family and finds important results about the state-dependence of care provision over time. In Hiedemann, Sovinsky, and Stern (2018), because the focus is on a sophisticated treatment of unobserved heterogeneity in a dynamic setting, the authors abstract from within-family strategic interaction. Instead, family members are modelled as having the same preferences over care outcomes, so the family can be treated as a unit. The key contribution of the current paper relative to Hiedemann, Sovinsky, and Stern (2018) is that children in the model are separate decisionmakers with their own preferences and interact non-cooperatively, allowing me to quantify the importance of strategic interaction over time.

[1]For instance, suppose Adam lives close to his mother but his sister Beth lives far away, and that their mother suddenly falls ill. If Adam and Beth cannot change location, then both Adam and Beth know that the probability of Beth providing care is very low, so Adam is effectively making decisions in isolation over whether to care for his mother. If, on the other hand, Beth is able to move back, or Adam is able to move away, then Adam will see the probability of Beth providing care as higher, so might be more tempted to shirk in the hope that Beth moves back and provides care.

Finally, with its emphasis on endogenous location choice and strategic shirking of a caregiving role the paper builds on work by Konrad et al. (2002), Maruyama and Johar (2017), Rainer and Siedler (2009), and Stern (2023). These papers all examine one-shot location decisions by adult children early in life seeking to avoid the burden of providing care to elderly parents in future[2]. The contribution of this paper relative to these others is that it allows for adult children to change location multiple times over their lives, meaning they can respond to parental health shocks. As outlined above, it is important to model this channel of behaviour in order to better capture the scope for strategic interaction between siblings over caring.

The rest of this paper is organised as follows. In Section 1.2 I set out some descriptive evidence about caregiving by adult children of elderly parents. In Section 1.3 I write down a dynamic model of interaction between siblings in the provision of care. Section 1.4 discusses data and estimation, providing intuitive arguments for identification and setting out results and model fit. Section 1.5 shows the results of counterfactual experiments and Section 1.6 concludes.

1.2 Descriptive evidence

I use data from the Health and Retirement Study (HRS) between 1998-2019[3]. The HRS asks respondents a battery of questions about their health, financial situation and relationships. A key advantage of the HRS, relative to other surveys of the older population in different countries, is its extensive use of proxies and following of respondents into nursing homes, meaning that it is more successful at mitigating sample attrition for respondents who develop care needs (Sonnega and Weir 2014; Weir, Faul, and Langa 2011).

In addition, HRS respondents are asked to provide information on any children they have. Notably, the HRS records for each child their income (in brackets), their education, their marital status, whether they live within 10 miles of the parent and how much care they provide to the parent. This allows me to link each parent with their possible child caregivers.

For calculating descriptive statistics and establishing general empirical patterns I use the full sample of respondents to these waves of the HRS who have at least one child record associated with them. This comprises of 35,105 unique respondents over the 11 biennial waves, with 200,385 respondent-wave observations, meaning that each respondent appears for an average of 5.7 waves. In Appendix 1.A.1 I show some descriptive statistics on the HRS sample.

[2]In Konrad et al. (2002), children make a one-shot location decision early in life and then parents can make a decision when they are elderly and need care to live close to a particular child. As discussed in Appendix 1.B, in my data it is children not parents who relocate when the parents need care.

[3]I use the RAND HRS Longitudinal File and the RAND HRS Family Data File and combine these with data from the RAND HRS Fat Files for each of the relevant waves (*HRS* 2024; *RAND* 2024).

I now use the HRS sample to present some key facts about the provision of care: namely, that there is a gender gap in the provision of care, the size of this gender gap depends on family composition, and most children do not receive substantial compensation for providing care.

1.2.1 There is a gender gap in the provision of care

Children are an important source of care to parents. Defining a parent as having care needs if they report difficulty with any Activity of Daily Living (ADL) or Instrumental Activity of Daily Living (IADL)[4], in the HRS sample I consider the average parent with care needs receives 20 hours of care per month from any daughters they have and 7 hours of care per month from any sons they have, together amounting to around a third of the 77 care hours per month that a parent with care needs receives from all sources[5]. For parents without a spouse, these figures are 32 hours and 12 hours respectively, making up 58% out of a total of 76 hours. In both cases, therefore, daughters provide around 3 times the care hours of sons. Appendix 1.A.2 discusses in more detail the breakdown of care receipt by source.

To identify some predictors of provision of care I use the child records associated with each respondent in the HRS. I run an OLS regression of a dummy for whether a given child provided any care to a parent with care needs on a set of child and parent demographics[6]. I also regress the log of hours of care provided on the same set of explanatory variables for the subsample of children who provided positive hours of care. These regressions are carried out for the full sample of HRS respondents with children, and are displayed in Table 1.1 below. The corresponding table employing a logit regression rather than a linear probability model, with qualitatively similar results, is presented as Table 1.8 in Appendix 1.A.3.

In Table 1.1, Columns 1 and 2 capture the extensive margin of care provision. In each case I regress a dummy for care provision on a set of parent and child characteristics. The difference between Columns 1 and 2 is that in Column 2 I include fixed effects for parent interacted with wave. As such, I consider the role of each child's characteristics relative to his or her siblings in a given wave in determining whether that child provides

[4]ADLs are very basic activities required for independent life such as basic mobility or dressing oneself. IADLs are more complex activities like preparing food or managing money. See Cleveland Clinic (2024) for more details.

[5]As shown in Appendix 1.A.2, the average parent with care needs receives 20 hours of care from daughters, 7 hours from sons, 28 hours from a spouse, 10 hours from other family members and 11 hours from non-family sources (not including employees of institutions).

[6]Note that if two members of a couple respond to the HRS, then any children they have will appear twice in the regressions below, once for each parent. In other words, the observations in this regression are unique parent-child links, rather than unique children.

Table 1.1: Predictors of care provision and the amount of care provided

	Dependent variable:		
	I(Hours help p.m. > 0)		log(Hours help p.m)
	(1)	(2)	(3)
Kid is daughter	0.022***	0.022***	0.216***
	(0.002)	(0.003)	(0.053)
Parent is mother	0.007**		0.001
	(0.002)		(0.048)
Daughter x Mother	0.063***	0.069***	0.201***
	(0.003)	(0.004)	(0.060)
Kid is eldest	0.005*	0.004	−0.024
	(0.002)	(0.003)	(0.032)
Kid is youngest	0.014***	0.002	−0.026
	(0.002)	(0.004)	(0.031)
Kid lives $\leq$ 10 miles away	0.144***	0.163***	0.543***
	(0.002)	(0.003)	(0.029)
Kid has sister	−0.019***		−0.089**
	(0.002)		(0.032)
Parent in couple	−0.053***		−0.114***
	(0.002)		(0.030)
Kid in couple	−0.027***	−0.026***	−0.368***
	(0.002)	(0.003)	(0.029)
Kid went to college	0.008***	0.016***	−0.177***
	(0.002)	(0.003)	(0.027)
Parent $\times$ wave FEs	N	Y	N
Observations	128 878	128 917	15 086
Adjusted R^2	0.161	0.349	0.119
Mean dep. var.	0.094	0.094	3.030

Notes: estimation via OLS. Only children of parents with care needs (ADL or IADL >0) are included. Regression weights are HRS respondent-level weights. In all regressions, other controls are dummies for number of ADL and IADL difficulties, education of parent, number of total kids of the parent, total number of grandkids of the kid, quadratics in age for kid and parent, a dummy for parent being White Caucasian, log of parent wealth and homeownership status of kid. For the fixed effects regression, the parent-level controls drop out. Standard errors clustered at the household level. *p<0.05; **p<0.01; ***p<0.001.

care to a particular parent[7].

Several important empirical patterns emerge. Note that as I include dummies for the kid being a daughter and the parent being a mother, the omitted category is son-father pairs. The coefficient on "Parent is mother" is very slightly positive, suggesting that sons are slightly more likely to provide care to mothers than they are to fathers. However, the dummy for "Kid is daughter" is positive and significant, so daughters provide more care than sons to fathers, and the interaction between "Kid is daughter" and "Parent is mother" is large and significant, so daughters provide much more care than sons to mothers[8].

Youngest children are slightly more likely to provide care than either eldest children or children who are neither youngest nor eldest, though this is much less of a predictor of care provision than the child being a daughter.

Kids are unsurprisingly much more likely to provide care if they live close to their parent. However, it is unclear whether these are children who were living within 10 miles of their parent before the health shock occurred (e.g. they never left the local area) or whether these are children who move back to provide care[9]. If the latter, then this suggests the welfare costs of providing care are greater, because they involve the cost in time and money of relocating, possibly changing jobs, and so on. Also, children moving back to provide care would point to a greater role for strategic interaction because children have more control over whether they are in a position to provide care. I examine this question in Appendix 1.B, considering kids' location choices around their parent developing care needs, exploiting the panel nature of the data to conduct a descriptive event study. The results offer some evidence that daughters in particular do move closer to parents following a parental health shock: on average, daughters are 1.0 percentage points more likely to live within 10 miles of their parent in the (two-year) period where the parent is first observed with care needs relative to the period before the parent develops care needs, compared to a 0.6pp increase for sons, relative to baselines of 38.7% for daughters and 37.6% for sons.

[7]Because the fixed effects capture any variation common to all children of a particular parent in a given wave - e.g. the gender or marital status of the parent - these explanatory variables drop out of Column 2. I also drop "Kid has sister" as a regressor for the fixed effect regression because I am already controlling for whether a given child's gender relative to their siblings through a combination of controlling for the child's gender and the parent $\times$ wave fixed effect.

[8]This interaction between parent and child gender suggests that at least part of the aggregate care gap is driven by the fact that women are overrepresented among the elderly needing care because fathers tend to die younger than mothers. Indeed, in the estimation sample for Table 1.1, 61% of parents are mothers. As a rough calculation, the results in Table 1.1 suggest that if instead there was a 50:50 split of mothers and fathers in the sample then other things being equal daughters would be 5.4pp more likely to provide care in the aggregate rather than 6.0pp, after controlling for all other regressors. This is because $0.61 \times (0.063 + 0.022) + 0.39 \times (0.022) = 0.060$ but $0.5 \times (0.063 + 0.022) + 0.5 \times (0.022) = 0.054$.

[9]Hiedemann, Sovinsky, and Stern (2018) consider the potential endogeneity of location, finding some evidence that location is endogenous, although the authors argue that any ensuing bias will be small in magnitude.

Kids are less likely to provide care if they have at least one sister[10], a phenomenon to my knowledge first noted by Grigoryeva (2017). The fact that children provide less care if they have a sister hints at the role of strategic interaction between siblings, discussed in more detail in the next section.

Both the parent having a spouse and the kid having a spouse are negatively associated with the probability of the kid providing care. The former is presumably due in part to there being a better "outside option" of care receipt for parents who have spouses; the latter may be because partnered kids have more demands on their time, or have a different set of parents-in-law that they need to spend time with, and so on.

More educated kids are more likely to provide care. Interestingly, this result holds even in Column 2, i.e. even controlling for unobserved heterogeneity at the parent-wave level through a fixed effect. In other words, kids who are more educated than their siblings are more likely to provide care than their siblings. This is striking because this provides some prima facie evidence against the hypothesis that children with a lower opportunity cost of providing care in terms of foregone wages will provide more care. Instead, it seems that within a particular family it is the more educated, and thus likely the higher earning, child who ends up providing more care. This point will be revisited when discussing the results of the estimation of the structural model.

Column 3 captures the intensive margin: for those providing positive hours of care, I regress the log of hours of help provided per month on the same explanatory variables as the first column. I do not include FEs here because many parents have only one child providing positive care hours to them in a given wave.

Similar patterns emerge. Daughters provide more care, even conditional on provision of care. This is particularly true when considering care given to mothers. Kids provide much more care if they live close, and slightly less care if they have a sister. It is interesting that conditional on providing care, having a spouse means a kid provides much less care, around two-thirds of what they would provide if they were single[11]. This negative association is much stronger than the negative association between the parent having a spouse and hours of care provided by the kid. Also, although more educated kids are more likely to provide positive hours of care, they provide less care conditional on positive provision than less educated kids.

[10]Note that one of the controls in this regression, as set out in the note beneath Table 1.1, is the number of children, so this dummy does not just capture the effect of the kid in question having at least one sibling.

[11]To check whether this is purely because of a kid's spouse stepping in to provide care, in Table 1.7 Appendix 1.A.3 I present the equivalent of Table 1.1 but counting all care provided by the kid and their spouse, rather than just the kid. The qualitative conclusions are the same - in particular, partnered kids provide fewer care hours conditional on providing positive hours.

1.2.2 The gender care gap depends on family composition

The fact that kids provide less care when they have at least one sister hints that they might reduce their care effort if they know a sister will "step up" in their absence. Here I consider this point in more detail.

In principle one could examine the role of strategic interaction between children by making use of data on only children in the data, i.e. children without brothers or sisters. If the ratio of care hours by only daughters to only sons is less than the ratio of care hours of daughters to sons within multiple child families, then this suggests that agents change their behaviour depending on the presence of siblings. Table 1.9 in Appendix 1.A.3 makes this comparison and finds that there is indeed a substantial difference.

The problem with using this approach to identify the role of strategic interaction in driving the gender care gap is that the number of children that a family has is likely endogenous. Parents have a significant degree of control over how many children they have, and thus the number of children they have will likely be correlated with observed and unobserved heterogeneity at the level of the parent.

Instead, I exploit variation in the gender composition of children, holding fixed the total number of children. In particular, I consider care provision by sons and daughters in two-child families. While gender composition could itself be endogenous[12] this is plausibly much less of a concern than for the number of children.

Figure 1.1 shows the probability of providing any care to a single[13] parent with care needs for different children in two-child families. The first bar is the probability of providing care by a daughter who has a sister, the second bar is probability of providing care by a daughter with a brother, and so on.

The figure suggests that daughters provide more care when their other sibling is a brother than when their other sibling is a sister; similarly, sons provide more care when their other sibling is a brother than when their other sibling is a sister. In other words, the graph suggests that relative to the case when their sibling is of the same gender as them, daughters increase their care and sons decrease their care when the sibling is of the opposite gender. An interpretation of this would be that sons shirk caregiving when they know their sisters will "step up" to fill the breach; conversely, daughters put in more effort to compensate for their brothers' shirking.

In Figure 1.8 in Appendix 1.A.3, I show the equivalent of Figure 1.1 while controlling for the location (whether living within 10 miles of the parent) of both children - similar

[12]For instance, if parents decide to keep having children until they have a daughter, then they are in some sense choosing the final gender composition of their children. Angrist and Evans (1998) use the sex mix of the first two children to predict whether a woman will have a third child, exploiting parents' preferences in general for having children of both sexes. Preferences for having children of both sexes - or alternatively preferences for having at least one daughter - might be plausibly correlated with e.g. expected care need in old age.

[13]I focus here on single parents because in married couples the partner tends to be the primary caregiver and because children of single parents will be the focus of my model.

Figure 1.1: Probability of providing care by type of kid

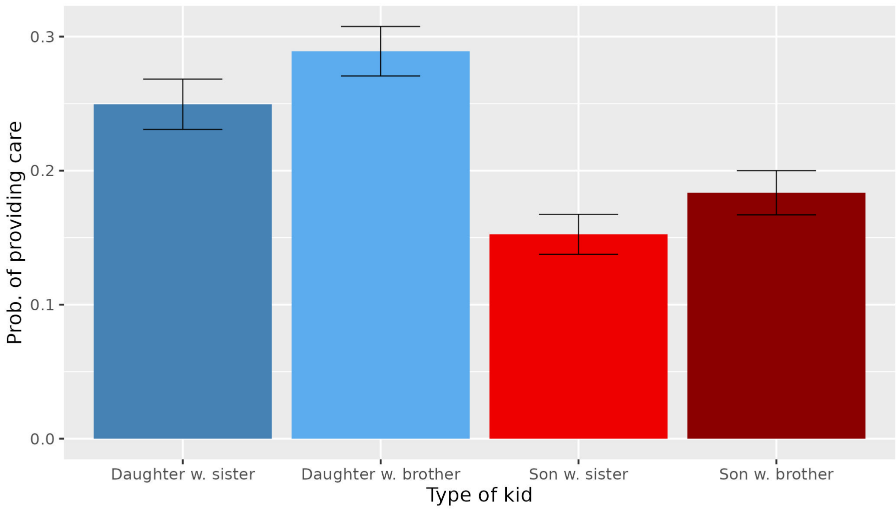

Notes: probabilities of providing care to a single parent with care needs for daughters whose only sibling is a sister, daughters whose only sibling is a brother, sons whose only sibling is a sister and sons whose only sibling is a brother. Means weighted by HRS respondent-level weights.

patterns persist. Also, in Figure 1.9 in the same Appendix, I present further information on the implications of this for total child-provided care hours received by the parent, and I discuss the issue of the influence of family composition of a person's family-in-law on their caring decisions. I argue that it is not of first-order importance to consider interaction between nuclear families as well as within nuclear families when it comes to the provision of care[14].

1.2.3 Most child caregivers do not receive substantial compensation

There is little prima facie evidence of child caregivers in the US receiving contemporaneous cash compensation from their parents (Barczyk and Kredler 2018; Groneck 2017). It is possible, however, that children are compensated by their parents in other ways. I consider here whether caregiving children receive higher bequests. In Appendix 1.C, I discuss two other candidate mechanisms of compensation, namely that children are compensated through the provision of rent-free accommodation and through the provision of childcare.

[14]More broadly, parents who only have sons see a drop off in the total hours of care they receive from their family, so daughters-in-law are clearly not generally stepping in and compensating for their husbands lack of care. For sick single parents with two children, those with two sons receive 39 hours of care from all family sources per month on average, compared to 54 hours for parents with one son, one daughter and 58 hours for parents with two daughters.

1.2.3.1 Compensation through bequests

Groneck (2017) uses 2002-2012 HRS Exit data on the bequests of single parents to argue that children who provide care receive more bequests than their non-caring siblings. He finds that, using family fixed effects to control for heterogeneity at the level of the family, a child who provides care is 5.4pp more likely to receive a bequest and, considering only children of parents with a positive estate to bequeath, children who provide care receive on average \$19.6k more than their siblings who do not provide care[15].

To assess this mechanism, I carry out a set of fixed effects regressions of measures of bequests on child characteristics, using 2002-2018 HRS Exit data, i.e. using three more waves of HRS Exit data[16]. In each case, the regression takes the form:

$$Y_i = \alpha_f + X_i\beta + \gamma CARE_i + \epsilon_i \tag{1.1}$$

where Y_i is the outcome variable[17], α_f is a fixed effect for family f, X_i is a vector of controls (see notes of Table 1.2), and $CARE_i$ is a dummy for whether child i provided care with ADLs or IADLs before the parent's death.

The main results of these regressions confirm the findings of Groneck (2017): using the longer HRS 2002-2018 exit sample, I find that a child who provides care is 4.9pp more likely to receive a bequest and considering only children of parents with a positive estate to bequeath, children who provide care receive on average \$14.3k more than their siblings who do not provide care[18].

However, as pointed out in Groneck (2017), many parents do not have an estate to bequeath in the first place: in my 2002-18 HRS Exit sample, only 54% of children have parents who die with a positive estate to bequeath. As such, it is difficult to infer from these headline results how much extra the *typical* child would expect to receive relative to their siblings if they provide care[19]. To assess this, I carry out a series of further regressions of the size of a bequest that a child receives on the same RHS variables but in each case changing the subsample of observations used.

Column 1 shows the baseline result, using only the children of parents with a positive estate to bequeath. In Column 2, I run the regression for all children, including the 46%

[15] Dollar values here and throughout, except where stated otherwise, are expressed in 2012 dollars using the CPI. As well as the family fixed effects model, Groneck (2017) also presents results from a 2SRI model and a Tobit model.

[16] As far as possible I use the same methods of sample selection and of construction of measures of bequests as in Groneck (2017). I would like to thank Max Groneck for sharing his replication code.

[17] This is either a dummy for receiving any bequest or the size of the bequest that child i receives in 2012 dollars.

[18] The regression table for the regression in terms of the amounts of bequests is given in Table 1.2. The regression table for the extensive margin of bequests, Table 1.11, is to be found in Appendix 1.C.1.

[19] Groneck (2017) does examine how bequests change along the wealth distribution by carrying out the regression in Equation 1.1 individually for each quintile of total family bequests, finding that in each case caregiving children receive higher bequests than their non-caregiving siblings. However, the regressions still condition on the parent having a positive amount to bequeath.

Table 1.2: Bequest received by care provided

	Dependent variable:				
	Bequest received ($1000 in USD2012)				
	(1)	(2)	(3)	(4)	(5)
$CARE_i$	14.283***	8.318***	6.731***	4.082***	3.118***
	(3.052)	(1.777)	(1.438)	(0.744)	(0.526)
Parent FEs	Y	Y	Y	Y	Y
Observations	6681	12 428	12 303	11 805	11 185
Adjusted R^2	0.828	0.838	0.851	0.802	0.730
Mean dep. Var	64.094	34.456	25.693	13.801	8.288

Notes: estimation via OLS. Column 1 considers only those children of parents with a positive estate at death. Column 2 considers all children. Column 3 drops children of parents with estates above the 99th percentile. Column 4 drops children of parents with estates above the 95th percentile. Column 5 drops children of parents above the 90th percentile. Controls are age of child, child education, child income, whether the child owns a home, whether the child lives within 10 miles of the parent, whether the child is co-resident with the parent and frequency of contact with the parent. Regression weights are HRS respondent-level weights.*p<0.05; **p<0.01; ***p<0.001.

of children whose parents die without a positive estate to bequeath. The coefficient on $CARE_i$ is cut almost in half as a result of expanding the sample.

The baseline result is also significantly driven by extreme results towards the top of the distribution of parent estates at death. In Columns 3, 4 and 5, I restrict the sample of Column 2. In Column 3, I drop children of parents with estates above the 99th percentile in size; in Column 4, I drop children of parents above the 95th percentile; and in Column 5, I drop children of parents with estates above the 90th percentile.

As more and more of the children of richer parents are dropped, the coefficient on $CARE_i$ unsurprisingly gets smaller. What is notable, however, is that even considering the bottom 90% of children in terms of parental estates (i.e. looking at Column 5), the coefficient is much smaller (around 5×) than the headline result[20]. A simple comparison of medians supports this tendency: the conditional median bequest received by child i relative to i's siblings (i.e. demeaned by family-specific means) is 0 regardless of whether $CARE_i$ is 0 or 1[21].

A concern with this analysis is the possibility of alternative causal links between caring

[20] According to Genworth (2024), formal care in the form of a full-time homemaker health aide costs $4760 per month in 2012 USD, suggesting that if we interpret the extra bequest for caregiving children as payment for their lifetime labour then these children are being paid a very small fraction of the market rate for care work, given that children often provide care for months or years at a time.

[21] In Figure 1.11 in Appendix 1.C.2 I compare all deciles of the bequest distribution for those with $CARE_i = 0$ and 1.

and receiving a bequest: for instance, providing care to a parent might prevent that parent resorting to formal care and thus running down a potential bequest, a mechanism that is part of the structural model in Section 1.3. In addition, children might be more motivated to provide care in cases where the potential bequest is larger. This would tend to bias the estimates in Table 1.2 upwards as the total bequest available, hence the maximum within-family difference in bequest between different children, is bigger when care has been provided by a child. Alternatively, parents with care needs likely face other important costs against which they cannot be fully protected by the provision of informal care, so these families could have smaller bequest pots on average at the time of death, biasing down the estimates in Table 1.2.

Reassuringly, the data suggest that controlling for the age of the parent at death and the number of children they have, receipt of care from a child is associated with a higher total bequest (across all recipients) on the part of the parent[22]. This indicates that the positive bias dominates so if anything the estimates presented in Table 1.2 can be seen as upper bounds on the true causal effect of a child providing care on their bequest receipt. This also provides motivation for the mechanism discussed in Section 1.3.2.4 whereby children can protect their bequest by providing care.

As such, it is interesting and important that on average, the mean caregiving child receives a significantly larger bequest than the mean non-caregiving child within a particular family. However, the results above suggest that the typical e.g. median caregiving child can expect much less in the way of compensation for providing care.

In Appendix 1.C I discuss two other mechanisms by which children might be compensated for providing care: through rent-free accommodation and through childcare. In neither case do I find convincing evidence that caregiving children receive substantial compensation through these channels. As such, overall, the evidence of this section suggests that although some child caregivers to parents may receive compensation for doing so, whether in the form of increased bequests, rent-free accommodation or childcare, this is not the case for most child caregivers. This is important because it implies the gender care gap is a welfare-relevant gap: it is not that daughters are more likely to provide care yet receive compensation for doing so, rather they provide largely uncompensated care. Also, this justifies the use of a model of non-cooperative interaction between siblings, insofar as there is little evidence of substantial within-family transfers alongside care arrangements, as would be expected in a cooperative model of interaction between siblings.

The results above suggest that many elderly people rely on their children to provide care, yet there is unequal sharing of the caring responsibility among the children. In

[22]See Table 1.12 in Appendix 1.C.3 for details.

particular, daughters are more likely to provide care. Sons who have a sister act strategically by reducing their care effort knowing their sisters will "step up" in their absence; conversely, daughters with brothers put in more effort to compensate for the fact that their brothers are less likely to provide care.

1.3 Model

1.3.1 Model overview

The model consists of a repeated game between two agents: Child A (elder) and Child B (younger). The two agents have a common single elderly parent. The parent is entirely passive in the sense that their decisions are entirely determined by state variables and their children's decisions.

At the start of each period, the parent experiences health and wealth shocks. After these shocks are realised, if the parent is still alive, the two agents simultaneously make (discrete) location, labour and care provision decisions and flow payoffs are realised.

If the parent has care needs and child i provides care but child j does not then both children derive a benefit from care being provided but only child i bears the cost of the provision of care. In this sense, the game resembles a public good provision problem: it is costly to provide the public good (informal care to a parent), yet both children derive a benefit if this good is provided. If neither child provides care then the parent pays for formal care out of pocket, decreasing the potential bequest that the children receive.

When the parent dies, then their remaining wealth is split equally between the children as a bequest. In other words, caregiving children do not receive more than non-caregiving children, in line with the lack of clear evidence for bequests being used to compensate children for caregiving.

1.3.2 Environment

Time is discrete, with each period lasting 2 years, matching the gap between HRS interview waves. The parent's health in period t is given by $h_t \in \{0, 1, 2, 3\}$, where $h_t = 0$ denotes that the parent is healthy, $h_t = 1$ denotes that the parent has moderate care needs, $h_t = 2$ denotes that the parent has severe care needs and $h_t = 3$ denotes that the parent is dead. At the start of the game, the parent is healthy, has age age_0 (with a maximum age of 100), and the game ends when the parent dies. The parent's health follows a first-order Markov process, with transition probabilities varying by current health, age and permanent income. I assume health transitions are exogenous and do not depend on informal care receipt, following the findings of Byrne et al. (2009) that care receipt does not substantively change health.

1.3.2.1 Choices

Each period the two children simultaneously make location, labour market and care decisions. Each child i can choose *Near* or *Far* as their location, where *Near* means the child lives within 10 miles of the parent and *Far* means the child lives further than 10 miles away[23], and each child can choose *No work* or *Work* as their labour market choice. If their parent is healthy ($h_t = 0$), children choose *No care provision* by default; if their parent has care needs ($h_t = 1$ or $h_t = 2$) children choose either *Care provision* or *No care provision*. Thus in general in each period t each child i makes a single discrete choice d_t^i out of the feasible set of choices F_t, where F_t has $2 \times 2 \times 1 = 4$ elements if the parent is healthy and $2 \times 2 \times 2 = 8$ elements if the parent has care needs.

1.3.2.2 Preferences

Child i's per period flow utility when the parent is alive is:

$$u_i(d_t^i, d_t^j, s_t, \epsilon_t^i) = g(c_t^i, l_t^i) + \omega(d_t^i, d_t^j, s_t) + \phi(d_t^i, s_t) + \epsilon_t^i(d_t^i) \tag{1.2}$$

The first component of utility is $g(\cdot)$, which captures utility from consumption c_t^i and leisure l_t^i. The other components of utility are $\omega(\cdot)$, which captures the (dis-)utility associated with the provision and receipt of care, $\phi(\cdot)$ which captures the (dis-)utility associated with the child's location choice, and a Type 1 Extreme Value preference shock ϵ_t^i associated with child i's discrete choice d_t^i. The (dis-)utility associated with the care choice depends on the set of common states s_t as well as both children's care decisions. The preference shock ϵ_t^i is iid across choices, children and time and has scale one.

The $\omega(\cdot)$ function takes the form:

$$\omega(d_t^i, d_t^j, s_t) = \omega_{pg}(d_t^i, d_t^j, s_t) + \omega_{warm}(d_t^i, s_t) \tag{1.3}$$

Here, $\omega_{pg}(\cdot)$ ("public good") captures the benefit to child i of their parent receiving informal care when sick, regardless of who provides the care. In other words, this is the benefit the child derives from the public good of having a cared-for parent[24], which is why it depends on child j's care choice as well as child i's.

The other term in the equation for $\omega(\cdot)$, namely $\omega_{warm}(\cdot)$, captures the "net warm glow" to child i of they themselves being the child who provides care - in other words, the net utility benefit associated with the act of providing care. This is part of utility does not depend on the other child's care choice. Note that it is plausible that the "net warm glow" of caring for a parent is negative, once the mental and physical burden of

[23]I assume all moving is done by children, and parental location is held fixed. As set out in Appendix 1.B, in my HRS descriptives sample approximately four times out of five when a child lives far in one wave and near in the next it is because the child has moved rather than the parent moving.

[24]In particular a parent who receives care from a family member, rather than from a formal source.

providing the care is taken into account. I defer the exact parameterization of $\omega_{pg}(\cdot)$, and $\omega_{warm}(\cdot)$ to Section 1.3.4.

1.3.2.3 Resources

As in Ko (2022), children divide their time between work, care time and leisure, and do not save any of their income. As such, they face two budget constraints, one for consumption and one for leisure:

$$c_t^i = w(d_t^i, s_t) \tag{1.4}$$

$$l_t^i = \bar{H} - H_{kt}^i - H_{wt}^i \tag{1.5}$$

In the first constraint, $w(\cdot)$ is the equivalised income function, determining child i's equivalised income (and hence consumption) as a function of their discrete choice d_t^i and the set of common states s_t (which will include elements like demographics). In the second constraint, $\bar{H}$ is the total amount of hours available per period, H_{kt}^i is hours spent caring and H_{wt}^i is hours spent working[25].

1.3.2.4 Parental wealth and bequests

As set out above, the parent's decisions in the model are entirely determined by state variables and their children's choices.

The parent starts the model with assets a_0. Their budget constraint takes the form:

$$a_{t+1} = a_t - (c_t^P - y^P) - ltc_t + b_t + \xi_t \tag{1.6}$$

Thus, assets tomorrow are equal to assets today minus consumption net of permanent income $(c_t^P - y^P)$, minus gross long-term care costs ltc_t, plus government benefits (including Medicaid) b_t and an iid wealth shock ξ_t, with distribution $N(\mu_\xi, \sigma_\xi^2)$.

I assume the parent always consumes their permanent income, so $c_t^P - y^P = 0$. Long-term care costs are 0 when the parent is healthy. I assume that government benefits are such as to guarantee each parent their permanent income in consumption[26], i.e.

$$b_t = max(0, ltc_t - a_t - \xi_t) \tag{1.7}$$

Thus, when the parent is healthy in period t, we will have that

[25] The mapping from discrete care/work choices to hours of care/work is set out in Table 1.3.

[26] This is not a substantive assumption: the decisionmakers in the model, i.e. the children, only care about whether their parent receives care and how much wealth they would receive as a bequest. Thus, having a Medicaid consumption floor and having parents consume at this lower level when they have no assets would produce the same behaviour in children, because parents' wealth would be the same.

$$a_{t+1} = max(0, a_t + \xi_t) \tag{1.8}$$

i.e. assets will follow a random walk (with drift, if $\mu_\xi \neq 0$), subject to the constraint that they are positive. Similarly, if the parent has care needs ($h_t > 0$) but receives care from at least one child, then they have no need to pay for formal care and the process for wealth is the same.

If, instead, the parent has care needs and receives no care from a child, then they face positive long-term care costs. The amount of out-of-pocket cost that a parent faces will depend on the severity of their condition. The process for wealth in this case will be:

$$a_{t+1} = max(0, a_t - ltc_t + \xi_t) \tag{1.9}$$

When the parent dies, their assets are split equally among their two children (regardless of who if anyone provided more care). Following Ko (2022), I assume that the children's terminal value is then the utility they would get from working full-time for the next T_{beq} periods, optimally splitting the consumption of the bequest over those periods, which provides a terminal payoff for the children to close the model[27].

In summary, apart from the direct utility benefit and cost to them of doing so, children are incentivised to provide informal care when their parent has care needs to stop their potential bequest being run down. This matches the mechanism in other papers like Ko (2022) and Barczyk and Kredler (2018).

1.3.3 Equilibrium

Let $\sigma^i(s_t, \epsilon_t^i)$ be child i's strategy, i.e. a mapping from the set of common states at t, denoted by s_t, and the vector of preference shocks, to the child's set of feasible actions in t, denoted by F_t, which will depend on the health of the parent[28]. Then, let $\sigma = \{\sigma^i, \sigma^j\}$ be the strategy profile across the children.

The child i's value function is:

[27]It may seem contradictory to include a bequest channel in the model when in Section 1.2 I argued that most children do not receive substantive compensation for their care in the form of bequests. However, the difference is that in Section 1.2 I was arguing that caregiving children do not tend to receive substantively more than their non-caregiving siblings. This is consistent with children still being motivated to provide care to prevent the overall level of the bequest - and hence their (equal) share of the bequest - from decreasing in size. Figure 1.4 in Section 1.4.6 shows that the model is able to match the association in the data between having low wealth and having care needs, suggesting that the model is able to capture the degree to which bequests can be run down by care needs.

[28]For instance, when the parent has care needs, F_t will equal {(0, 0, 0), (0, 0, 1), (0, 1, 0), (0, 1, 1), (1, 0, 0), (1, 0, 1), (1, 1, 0), (1, 1, 1)}, where the first element in each triple is the care choice, the second element is the location choice and the third element is the labour market choice; when the parent is healthy, F_t will consist only of the first four of these eight triples.

$$V^i(s_t, \epsilon_t^i, \sigma) = \max_{d_t^i \in F_t} \left\{ \mathbb{E}[u^i(d_t^i, d_t^j, s_t, \epsilon_t^i) + \beta V^i(s_{t+1}, \epsilon_{t+1}^i, \sigma) | s_t, d_t^i, \sigma] \right\} \tag{1.10}$$

Let the child i's choice-specific value function for d_t^i be their expected flow payoff from choosing d_t^i, less the value of the preference shock associated with d_t^i, plus the expectation of their discounted future value function:

$$v^i(d_t^i, s_t, \sigma) = \mathbb{E}[\pi^i(d_t^i, d_t^j, s_t) + \beta V^i(s_{t+1}, \epsilon_{t+1}^i, \sigma) | s_t, d_t^i, \sigma] \tag{1.11}$$

where $\pi^i(d_t^i, d_t^j, s_t) = u^i(d_t^i, d_t^j, s_t, \epsilon_t^i) - \epsilon_t^i(d_t^i)$.

Then, the strategy profile σ^* is a Markov Perfect Equilibrium iff:

$$\sigma^{*i}(s_t, \epsilon_t^i) = \text{argmax}_{d_t^i \in F_t} \left\{ v^i(d_t^i, s_t, \sigma^*) + \epsilon_t^i(d_t^i) \right\} \tag{1.12}$$

for all t, for all s_t, for all ϵ_t^i and for all $i \in \{A, B\}$. This equilibrium condition captures the idea that in an equilibrium, each player's strategy is a best response to all other players' strategies.

The model will not in general have a unique equilibrium, because players move simultaneously when making decisions. I will assume that in each family in the data the same unique equilibrium is being played. This is a strong assumption but it is difficult to allow for multiple equilibria in estimation given the relatively small number of periods when I observe each family[29]. Alternatively, I could change the model to ensure there is a unique equilibrium by specifying that players move sequentially in each period, with knowledge of the previous player's move, but it is not clear on what basis to select the first-mover in each case, and the model would be assuming significant differences between the first- and second-moving child that plausibly do not exist in reality. Moreover, the estimation results in Section 1.4 indicate that children's care decisions are not strongly correlated conditional on observables suggesting that there is not a significant problem of multiple equilibria or other unobserved heterogeneity at the family level in the data. I do however account for the possibility of multiple equilibria when conducting counterfactual exercises in Section 1.5.

1.3.4 Functional forms and parameterisation

In this subsection I discuss the exact functional forms I choose to provide structure to the model.

[29]This stands in contrast to examples from the industrial organisation literature, where long panels for each market mean that an econometrician could allow for each market to have its own equilibrium, see e.g. Aguirregabiria, Collard-Wexler, and Ryan (2021).

1.3.4.1 Utility over consumption and leisure - $g(\cdot)$

I assume that agents have Cobb-Douglas preferences over consumption and leisure:

$$g(c_t^i, l_t^i) = \theta_c \log(c_t^i) + \theta_l \log(l_t^i) \tag{1.13}$$

The scale of the parameters is normalised by the scale parameter of the iid Type 1 Extreme Value preference shock in Equation 1.2, which I set equal to 1.

1.3.4.2 "Public good" component of care utility - $\omega_{pg}(\cdot)$

As set out above, $\omega_{pg}(\cdot)$ captures the benefit to a child of their parent receiving some form of informal care, regardless of which child provides it. Letting k_t^i be a dummy for whether child i provides care at t, I write it as:

$$\omega_{pg}(\cdot) = \begin{cases} \alpha_{h1}, & \text{if } h_t = 1 \ \& \ k_t^i + k_t^j > 0. \\ \alpha_{h2}, & \text{if } h_t = 2 \ \& \ k_t^i + k_t^j > 0. \\ 0, & \text{otherwise.} \end{cases} \tag{1.14}$$

Thus, for instance, if the parent has health $h_t = 1$ and at least one child provides care ($k_t^i + k_t^j > 0$) then both children receive benefit α_{h1}, regardless of who actually provided the care. The benefit they derive potentially varies with the severity of the health condition of the parent.

1.3.4.3 "Net warm glow" component of care utility - $\omega_{warm}(\cdot)$

The function ω_{warm} captures the utility net benefit to child i of they themselves providing care (i.e. regardless of what their sibling does):

$$\omega_{warm}(\cdot) = \begin{cases} X_t^i \gamma & \text{if } h_t > 0 \ \& \ k_t^i > 0. \\ \\ 0, & \text{otherwise.} \end{cases} \tag{1.15}$$

where X_t^i is a set of characteristics of the child and parent. In particular, X_t^i contains controls for:

- The severity of parental health problems ($h_t = 1$ or $h_t = 2$). It is plausible both that the cost of providing care, and the "warm glow" from doing so (or guilt from not doing so) will vary with the severity of the parent's health problems;
- Whether the child lives further than 10 miles from the parent at the time of providing care. This is to capture the extra time and money spent in travelling to provide care;

- Whether the child lived further than 10 miles from the parent at the start of the game. Briefly, this is because it is possible that children who live far from their parents at the start of the game are systematically different from those who live near in terms of unobserved attitude or affection towards a parent. This issue is discussed in more detail in Appendix 1.D;

- Whether the child is a son. Allowing the net benefit of providing care to depend on child i's gender means that the model allows for there to be unobserved preference differences between men and women when it comes to the provision of care to their parents. In principle, this simple preference difference could capture lots of different types of differences in the situations of sons and daughters. For instance, if the coefficient on son^i is negative, this could capture the fact that daughters genuinely derive more enjoyment from providing care than sons; or that daughters are "better" at providing care than sons in the sense that for a given unit of costly effort daughters will provide more effective care than sons; or that parents prefer receiving care from daughters rather than sons, and it is costly to go against parents' wishes; or that it is costly to go against some norm in society regarding daughters' role as primary caregiver; or some combination of these explanations. For instance, Byrne et al. (2009) argue both that daughters are better at providing care than sons and experience less burden from doing so. For the sake of tractability of the model I do not distinguish between these separate drivers of preference difference between sons and daughters;

- Whether the parent is a father, and an interaction between parent and child gender. This is to allow for the stylised fact from Section 1.2 that mothers receive more care overall and that daughters seem to provide relatively more care to mothers while sons provide relatively more care to fathers;

- Whether the child is switching from non-provision to care provision this period. This is to capture an important dynamic aspect of care provision, namely the inertia of care provision arrangements: families tend to have a primary caregiver which does not change over time (Hiedemann, Sovinsky, and Stern 2018). It is important to include this term for the sake of capturing the nature of strategic interaction between children. If providing care in one period causes a child to be "locked in" to the caregiving role, in the sense that the cost of them providing care in subsequent periods will be much lower than that of their siblings, then this makes the initial strategic interaction between the children at the point of a parent developing care needs more significant in explaining the distribution of care roles into the future;

- Whether the child is the younger of the two children ("seniority"), and an interaction between seniority and child gender, to capture any social norms or preference differences connected with being the eldest/youngest child, and in particular with being an eldest/youngest daughter.

I write Equation 1.15 out in full in Appendix 1.D, where I also discuss the issue of unobserved heterogeneity in preferences for providing care.

1.3.4.4 Location choice utility - $\phi(\cdot)$

The flow utility from location choice is comparatively simple:

$$\phi(\cdot) = \phi_{far} far_t^i + \phi_{move} I(far_t^i \neq far_{t-1}^i) \tag{1.16}$$

Thus, children derive (dis-)utility from living far from their parents, and also derive (dis-)utility from changing their location. These parameters thus capture whether overall it is more attractive to live far from one's parents for whatever reason, and how costly it is to move.

1.4 Estimation

1.4.1 Estimation sample

For the estimation sample I use data from HRS Waves 4 to 14 (1998-2018). I select those families with one surviving parent and exactly two child records in a given year, where the parent is aged between 55 and 85 in the first observation in the data[30].

I drop any parents who hold long-term care insurance[31]. I drop any cases where in the first period where the family is observed the parent with care needs and either child provided care in the previous period. I do this so at the start of the sample I avoid an "initial conditions" problem of children already having selected into caregiving, as discussed in Appendix 1.D. Finally, in cases where data on care, work, location, education or marital status of child, or health of the parent, are missing for a single wave I impute this from the child's data in the previous wave but if the data are missing for more than one consecutive wave I drop the family from the first wave with missing data. I am left with an estimation sample of 8444 family-wave observations, consisting of 2459 families who are in the sample for an average of 3.4 waves each.

I estimate each parent's permanent income as the mean of their income for each period they are observed in the data, and let parental health transition probabilities vary by whether the parent has above or below median population permanent income.

[30] If the family is observed at parent ages before 55, I drop all observations until the parent is 55.

[31] I drop these parents because I do not model LTCI choice or coverage. Only 10% of the full HRS sample has LTCI.

To avoid having to solve the model separately for every single possible starting age of the parent, I instead separate parents into two groups according to their initial age, and treat this initial age group as an additional non-time-varying state variable. In other words, I solve the model separately for these two different initial age groups. Specifically, any parent in the 55-69 age group in their first observation I class as starting the model with age 63 and anyone in the 70-85 age group I class as starting the model with age 78. This simplifies the solution because it means that I only need to solve the model for two age groups rather than as many age groups as there are starting ages in the data.

1.4.2 Parameters from outside the model

Many of the parameters of the model I set using values from the literature or directly from the data. The choices are summarised in Table 1.3 below.

Table 1.3: External parameters

Parameter	Value	Source
Health transition probabilities	-	HRS
Formal care cost p.a., $h_t = 1$	\$15.0k	Genworth (2024)
Formal care cost p.a., $h_t = 2$	\$74.6k	Genworth (2024)
Informal care hours p.w., $h_t = 1$	13	HRS
Informal care hours p.w., $h_t = 2$	25	HRS
Full-time work hours p.w	35	Author's choice
μ_ξ	-\$2.0k	HRS
σ_ξ	\$115.3k	HRS
Income process	-	PSID
Initial age	{63, 78}	HRS
Terminal age	100	Author's choice
T_{beq}	5	Author's choice
Age gap to elder child	24	HRS
Age gap to younger child	29	HRS
β	0.93	Author's choice

Notes: see Appendix 1.E for detail on how these values are chosen or estimated.

In Appendix 1.E, I discuss in more detail how I arrive at these external parameters.

1.4.3 Parameters estimated inside the model

I estimate the rest of the model by maximising the pseudo likelihood of agents' choices (Aguirregabiria and Mira 2007). There are 16 parameters to be estimated, summarised in Table 1.4 below.

Table 1.4: Parameters estimated inside the model

Parameter	Description
Consumption and leisure	
θ_c	Weight on consumption
θ_l	Weight on leisure
Child's "public good" utility of providing care given that...	
α_{h1}	Parent is in moderate bad health
α_{h2}	Parent is in severe bad health
Child's "net warm glow" utility of providing care given that...	
γ_{h1}	Parent is in moderate bad health
γ_{h2}	Parent is in severe bad health
γ_{son}	Child is a son
γ_{dad}	Parent is a father
γ_{start}	Child starts providing care this period
$\gamma_{origfar}$	Child originally lived far from parent
γ_{far}	Child currently lives far from parent
$\gamma_{youngest}$	Child is youngest sibling
$\gamma_{youngest \times son}$	Child is youngest sibling and is a son
$\gamma_{dad \times son}$	Parent is a father and child is a son
Child's utility from location choice given that...	
ϕ_{far}	Child lives far from parent
ϕ_{move}	Location choice utility, child moved this period

Notes: see Section 1.3 and Appendix 1.D for discussion of how these parameters enter into a child's utility function.

I follow Ko (2022) in using the insight of Bajari, Benkard, and Levin (2007) regarding exploiting the linearity of the flow payoff function, thus the value function, to speed up estimation.

My approach is as follows. I start by estimating agents' conditional choice probabilities (CCPs) in each state using a logit regression of observed choices in the data on state variables[32]. I then use these CCPs to construct the value functions via simulation.

Specifically, let the true probability of child i choosing discrete choice d_t given state variables s_t and strategy profile σ be:

$$P^i(d_t|s_t,\sigma) = \int I(\sigma^i(s_t,\epsilon_t^i) = d_t) f(\epsilon_t^i) d\epsilon_t^i \tag{1.17}$$

i.e. the probability that d_t is the option "chosen" by i's strategy given the preference shocks.

I estimate the sample analogue of Equation 1.17 from the data using a logit regression, regressing observed choices on the set of observed common states in the data s_t. This produces estimates of the form $\hat{P}^i(d_t|s_t)$[33].

Using the fact that preference shocks are iid Type 1 Extreme Value, the choice-specific value function for choice d_t, relative to some reference choice d_t^0, will be:

$$v^i(s_t,d_t,\sigma) - v^i(s_t,d_t^0,\sigma) = \ln P^i(d_t|s_t,\sigma) - \ln P^i(d_t^0|s_t,\sigma) \tag{1.18}$$

and thus I can recover an estimate of the strategy of child i using:

$$\hat{\sigma}^i(s_t,\epsilon_t^i) = \text{argmax}_{d_t^i \in F_t} \left\{ \ln \hat{P}^i(d_t^i|s_t) - \ln \hat{P}^i(d_t^{i0}|s_t) + \epsilon_t^i(d_t^i) \right\} \tag{1.19}$$

which is derived by substituting the sample analogue of Equation 1.18 into the definition of equilibrium in Equation 1.12[34].

Given the policy function estimates $\hat{\sigma}$, I estimate the value functions by using $\hat{\sigma}$ to simulate forward and summing up flow payoff in each period. I do this for S simulation draws and then take the mean of each simulated value function as my estimate. As in Ko (2022), because the flow payoff functions, hence the value functions, are linear in parameters, this simulation procedure has only to be done once. This is because the value function $V(s_t,d_t,\sigma)$ can be written as $W(s_t,d_t,\sigma)\theta$, where θ is the vector of parameters to be estimated and where $W(s_t,d_t,\sigma)$ does not depend on unknown parameters, so I

[32] The state variables in the model are the gender, starting location, education, marital status, previous location, previous caring and previous labour force decision of each child; the gender, health, wealth, permanent income group, starting age (older/younger) and age difference from starting age of the parent.

[33] There is no equivalent of σ in this expression because strategies are not directly observed. Recall that I assume that there is only one equilibrium being played in the data, thus all children must be using the same strategy conditional on whether they are the older or younger child.

[34] To do this one must recognise that $\text{argmax}_{d_t^i \in F_t} \{v^i(d_t^i,s_t,\sigma) + \epsilon_t^i(d_t^i)\} = \text{argmax}_{d_t^i \in F_t} \{v^i(d_t^i,s_t,\sigma) - v^i(d_t^{i0},s_t,\sigma) + \epsilon_t^i(d_t^i)\}$ for reference choice d_t^{i0}.

only need to estimate $W(s_t, d_t, \sigma)$ once and then scale by the candidate parameter value θ. I then choose the θ to maximise the pseudo likelihood of the observed choices, as in Aguirregabiria and Mira (2007). This process is discussed in more detail in Appendix 1.F.

1.4.4 Identification

I here provide some informal arguments for how the parameters of the model can be identified by variation in the data.

The relative size of the consumption and leisure preference parameters is pinned down by agents' work and care choices. First, the more that people in general choose full-time work rather than no work, the more important is consumption relative to leisure for them. Also, working hours are the same regardless of demographics but certain demographics (e.g. being a man) are associated with a bigger income, hence consumption, when working relative to not working. This provides extra variation in the incentive to work rather than to take leisure which will provide better identification of the relative sizes of the consumption and leisure parameters.

The preference for living far from one's parent is pinned down by the proportion of children who live far from their parent, and the transition cost of changing location is pinned down by the rate of transition between the two location statuses.

The overall motivation to provide care – i.e. the sum of the "public good" and "net warm glow" parameters – is pinned down by different rates of providing care according to different demographics in the data.

Finally, the relative sizes of the "public good" and "net warm glow" parameters is pinned down by variation in the probability of each child's sibling providing care. Consider two children, Arnold and Bob, who belong to two different families. Arnold and Bob are identical apart from the fact that Arnold's sibling has demographics associated with providing a lot of care (e.g. Arnold has a sister) but Bob's sibling has demographics associated with not providing much care (e.g. Bob has a brother). If Bob is just as likely to provide care as Arnold, this suggests that the "public good" motivation to provide care is much less than the "net warm glow" motivation, because Bob's lower probability of having a sibling "step up" to cover Bob's failure to provide care does not change how likely he is to provide care. However, if Bob is more likely to provide care than Arnold, this suggests that the "public good" motivation is important. In other words, Bob is more motivated to provide care because he knows it is unlikely his sibling will provide the public good if he does not, whereas Arnold is comfortable shirking and letting his sister provide care. Thus, the extent to which people like Bob provide more care than people like Arnold pins down the relative roles of "public good" motivations and "net warm glow" motivations.

1.4.5 Results

The results of the estimation are given in Table 1.5 below. I estimate standard errors using 20 bootstrap replications.

For the Baseline model (Column 1), the key results are as follows. Children in the model are estimated to have low preference for leisure relative to consumption[35]: the estimates imply that on average agents would be willing to give up 16% of their leisure time for around a 10% increase in their equivalised consumption.

As for the "public good" care choice parameters, having a parent with either moderate or severe care needs receive informal care from their children is a public good for those children: both children derive a notable positive benefit regardless of who provides the care, though the coefficient for severe health problems is not significant at the 5% level (though it is significant at the 10% level). Interestingly, the public good benefit seems to be larger when the parent's health problems are moderate rather than severe, though the difference is not statistically significant.

As for the "net warm glow" care choice parameters, it is important to recognise that the reference category of child in the estimation is a daughter who lives close to their parent (both now and at the start of the game), who provided care the previous period, who is the elder of the two children and whose parent is a mother[36]. Thus, γ_{h1} and γ_{h2} reflect the net warm glow to this type of child from providing care when the parent has moderate and severe care needs respectively. The other γ parameters modify this net warm glow for other types of child.

Sons have a bigger utility cost of providing care than daughters for providing care. In particular, γ_{son} is negative and significant, which indicates that for sons who are the eldest child of a mother vs. daughters who are the eldest child of a mother, the sons suffer a bigger cost. While not statistically significant, the coefficient $\gamma_{younger \times son}$ is slightly positive, suggesting that this gap is slightly smaller when considering sons who are the youngest child vs. daughters who are the youngest child, and the coefficient on $\gamma_{dad \times son}$ is positive, suggesting that the gap is smaller when considering sons of fathers vs. daughters of fathers.

To provide a sense of the economic significance of the estimated parameters, the estimates imply that if an elder daughter of a mother was indifferent between providing care or not, an agent who was identical apart from the fact that he was a son instead would have to receive a 31% increase[37] in equivalised consumption as compensation for

[35] The scale of all coefficients is determined by the scale parameter of the Type 1 Extreme Value preference shock, which I have normalised to 1.

[36] This is because I estimate coefficients on dummies for being a son, for living far from the parents in the current period and at the start of the game, for starting providing care this period, for being the younger of the two children and for the gender of the parent, thus the omitted category is as described.

[37] This comes from $0.455/\ 1.478 \approx 0.31$, where 1.478 is the coefficient on the log of equivalised consumption.

Table 1.5: Estimation results

	(1) Baseline	(2) No PG utility	(3) No gender diff.	(4) No health diff.
Consumption and leisure params				
θ_l - weight on leisure for K	0.948***	0.950***	0.952***	0.939***
	(0.058)	(0.059)	(0.062)	(0.059)
θ_c - weight on consumption for K	1.478***	1.478***	1.479***	1.473***
	(0.034)	(0.035)	(0.034)	(0.033)
Public good care params				
α_{h1} - P in moderate bad health	1.539**	-	1.750***	1.115*
	(0.539)	-	(0.527)	(0.459)
α_{h2} - P in severe bad health	1.104	-	1.223*	-
	(0.580)	-	(0.557)	-
Net warm glow care params				
γ_{h1} - P in moderate bad health	-0.682	0.602***	-1.037*	-0.146
	(0.505)	(0.159)	(0.490)	(0.438)
γ_{h2} - P in severe bad health	0.011	0.884***	-0.261	-
	(0.524)	(0.158)	(0.491)	-
γ_{son} - K is a son	-0.455*	-0.508*	-	-0.461*
	(0.217)	(0.215)	-	(0.216)
γ_{start} - K just starting caring	-1.201***	-1.327***	-1.232***	-1.260***
	(0.141)	(0.122)	(0.139)	(0.139)
$\gamma_{origfar}$ - K originally lived far	0.037	0.011	0.053	0.055
	(0.128)	(0.126)	(0.125)	(0.124)
γ_{far} - K currently lives far	-0.919***	-0.954***	-0.917***	-0.927***
	(0.170)	(0.170)	(0.171)	(0.167)
$\gamma_{younger}$ - K younger of two siblings	0.107	0.124	0.130	0.112
	(0.139)	(0.131)	(0.118)	(0.135)
$\gamma_{younger \times son}$ - K younger & K son	0.045	0.031	-	0.033
	(0.295)	(0.289)	-	(0.292)
γ_{dad} - P is father	-0.434*	-0.343*	-0.278	-0.436*
	(0.182)	(0.163)	(0.159)	(0.174)
$\gamma_{dad \times son}$ - P father & K is son	0.372	0.387	-	0.419
	(0.276)	(0.270)	-	(0.270)
Location choice params				
ϕ_{far} - K currently lives far	0.041**	0.045**	0.040*	0.041**
	(0.016)	(0.016)	(0.016)	(0.016)
ϕ_{move} - K moved location	-1.494***	-1.495***	-1.494***	-1.494***
	(0.012)	(0.012)	(0.012)	(0.012)
LR test statistic	-	16.211	13.718	9.046
p-value	-	<0.001***	0.001**	0.011*

Notes: "P" refers to the parent and "K" refers to the kid. Estimation via PML maximization. Standard errors calculated using 20 bootstrap replications.*p<0.05; **p<0.01; ***p<0.001.

providing care to be indifferent between providing care and not.

Unsurprisingly there are major start-up costs associated with providing care: γ_{start} is negative and significant. This is consistent with the results in Ko (2022) and Skira (2015). Whether a child lived far from the parent at the start of the game does not have a significant impact on their net benefit of providing care, as $\gamma_{origfar}$ is estimated as a noisy zero. In other words, these children do not seem to be systematically different (above and beyond their original location choice) from other children who live near to the parent at the start of the game. However, whether or not a child lives near the parent at the point of providing care is very important to their net benefit of providing care: the coefficient γ_{far} is negative and significant.

Younger children do not seem to be systematically different to older children when it comes to the provision of care as neither $\gamma_{younger}$ nor $\gamma_{younger \times son}$ is significant.

The results for the gender of the parent match up with the stylised empirical facts. Fathers are less likely to receive care than mothers ($\gamma_{dad} < 0$), but the gender gap in care for fathers is less than the gender gap in care for mothers ($\gamma_{dad \times son} > 0$), though this latter coefficient is not statistically significant.

Finally, as for the location choice parameters, children have a slight preference for living far from their parent. Moreover, switching from living far to living near or vice versa incurs a major utility cost for the child as ϕ_{move} is negative and significant. This allows the model to match inertia in location choice over time.

1.4.5.1 Alternative specifications

Columns 2 to 4 of Table 1.5 present some alternative specifications of the model. In Column 2, the "public good" motive to provide care is set to 0 (i.e. $\alpha_{h1} = \alpha_{h2} = 0$), so agents derive utility from their parent receiving care only if they are the ones providing it. In Column 3, sons and daughters have the same preferences over providing care to parents (i.e. $\gamma_{son} = \gamma_{younger \times son} = \gamma_{dad \times son} = 0$). In Column 4, children's caregiving preferences are the same regardless of whether the parent is in moderate bad health or severe bad health (i.e. $\alpha_{h1} = \alpha_{h2}$ and $\gamma_{h1} = \gamma_{h2}$).

For Column 2, the one notable change to the estimated coefficients is that γ_{h1} and γ_{h2} are now positive and significant because children need an "extra" motivation to provide care given the "public good" motivation is no longer operative. The model fit is worse: in the final rows of the table, I report the result of a Likelihood Ratio test relative to the baseline model, where it can be seen that the test statistic is significant at the 0.1% level[38].

For Column 3, it is notable that the "public good" parameters are bigger and the "net warm glow" parameters are smaller (more negative). Intuitively, this is because when

[38] For the purposes of carrying out the test I treat the pseudo likelihood as if it were a likelihood, so I am only approximating the true LR test.

there are no preference differences between sons and daughters, there is less strategic shirking in the model, because there are fewer cases where one child thinks that their sibling is in a much better position than they are to provide care. Thus, to match the amount of strategic shirking in the data, more weight is given to the "public good" motivation to provide care as this increases strategic shirking. The model fit is much worse: the LR ratio test statistic is highly significant.

Finally, for Column 4, the only coefficients that change notably are α_{h1} and γ_{h1}, with both shifting towards 0. Again, the model fit is significantly worse.

1.4.6 Model fit

To assess model fit I generate simulated data on care, work and location choices for 10000 families using the model.

For each simulation, I draw a family's initial observation from the data, assign that family the health states of the parent observed in every period in the data, and simulate the choices they would make according to the estimated model. The family drops out of the simulation in the same period they drop out of the data.

Figure 1.2 shows the results of an OLS regression of a dummy for a child providing any care, conditional on their parent having health problems, on a set of observables, for both the real data and the model simulated data.

The fit is broadly good. In particular, the coefficients on the dummy for whether the child provided care the previous period, the dummy for the child being a daughter and the dummy for whether the child lived within 10 miles in the previous period match well across the real data and the simulated data. It is notable that the coefficient on "Sib cares", i.e. a dummy for whether the child's sibling provides care in the current period, is not statistically significant. If it were statistically significant that would suggest that the model had been mis-specified, because if preference shocks are iid over time and across players then conditional on observables agents' decisions should be independent (Aguirregabiria and Mira 2019).

There is one notable failure of model fit however: the coefficient on the dummy for whether the child went to college. This is positive and significant in the data but close to 0 in the model, so the model does not replicate the empirical pattern that more educated children are more likely to provide more care. The mechanical reason for this is that the only role of education in the model is in the income equation. As shown in Appendix 1.E, the coefficient on College $\times$ Works is actually slightly negative, suggesting that the gap between log income of college and non-college educated children is smaller when both types of children are working than when they are not working, or that in other words college educated children have a lower opportunity cost of not working, but this differential is not big enough to drive more educated children to provide more care than

Figure 1.2: Results from regression of care dummy on observables - data versus model

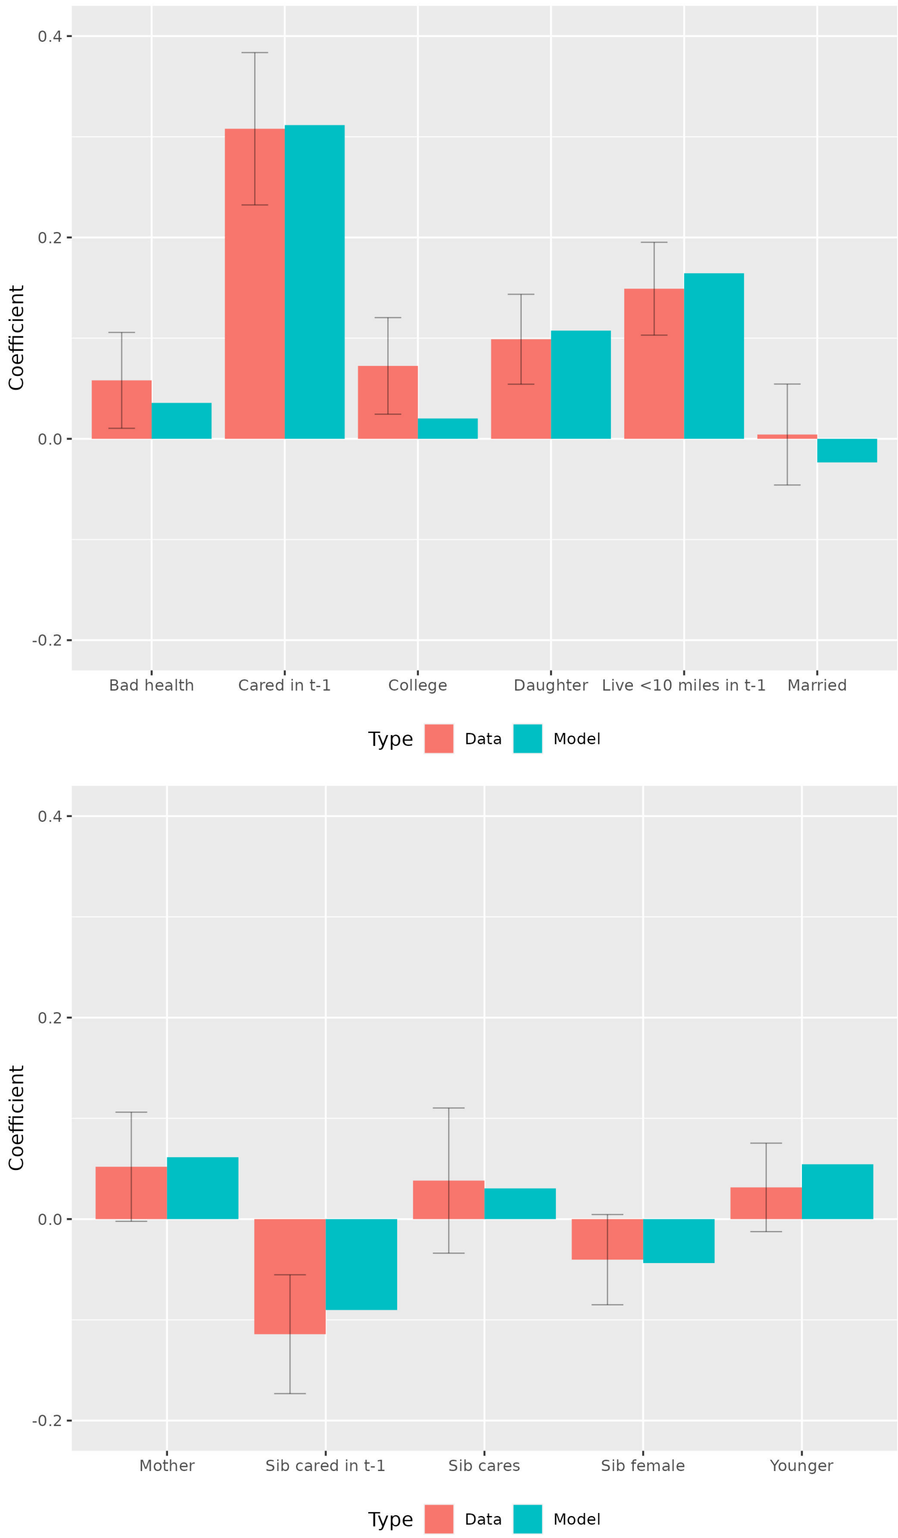

Notes: estimation via OLS. Standard errors clustered at household level. Error bars represent 95% confidence intervals. For the real data, N=1474. "Bad health" is a dummy for whether the parent is in the most severe health category. "Sib" refers to the child's sibling, so e.g. "Sib cares" is a dummy for whether the child's sibling provided care in the current period. Other regressors (not shown) are the education, marital status, location and care provided last period of the sibling, as well as the inverse hyperbolic sine of wealth and a polynomial in age; in no case were these extra omitted coefficients significantly different from their simulated counterparts at the 5% level.

less educated children. Thus assuming no preference difference for more educated children implies that the model will be unable to match the fact that more educated children provide more care. Why exactly education has a positive impact on care provision is an important and interesting question but is beyond the scope of this paper.

I also assess the model's fit when it comes to the work decision. To do this, I use the simulated data to calculate the proportion of people working in every period and compare these rates to those from the real data. The results are shown in Figure 1.3.

Figure 1.3: Proportion of children working - data versus model

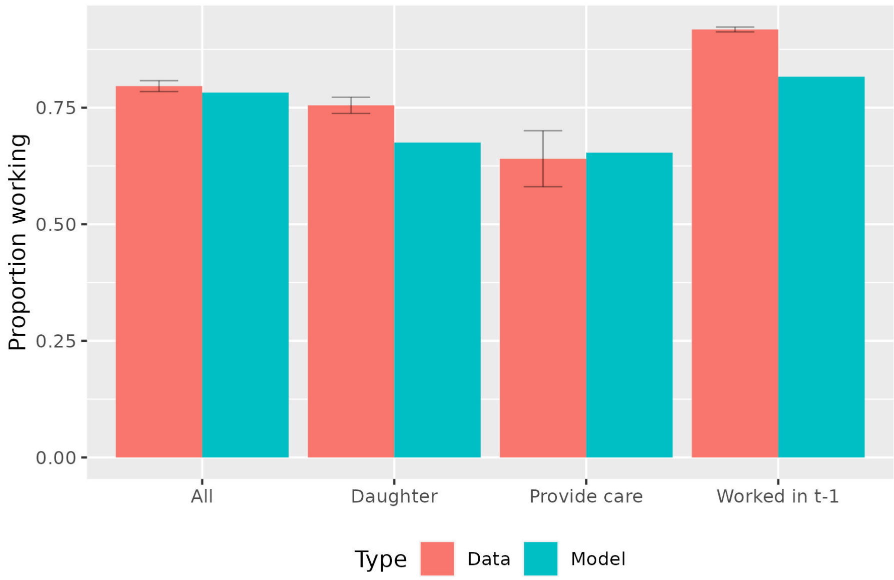

Notes: standard errors clustered at household level. For the four cases, N in the real data is 16888, 8593, 434 and 13642 respectively.

Again, the fit is reasonably good. The model matches the proportion of people working in the real data well, and also matches the proportion of people working while also providing care, and thus seems to capture the trade-off between working more and providing care to one's parent. However, the model understates the proportion of daughters who work (and correspondingly overstates the proportion of sons who work). Also, the model understates the persistence of working arrangements: in the data, conditional on working in $t-1$, the probability of working in t is 92%, whereas in the model it is only 82%. One explanation for this is that the only source of persistence in working arrangements in the model is that there is a wage penalty for not having worked the previous period, incentivising agents to not take breaks from working for a period. A more sophisticated model of labour market choices would allow for different preferences for work in the population so that those who are less work-averse select into working, adding a different form of persistence into the labour market choice.

Finally, Figure 1.4 below plots the proportions of parents with low wealth (below \$1000) in the model and the data, unconditionally and by health status.

Figure 1.4: Proportion of parents with low wealth - data versus model

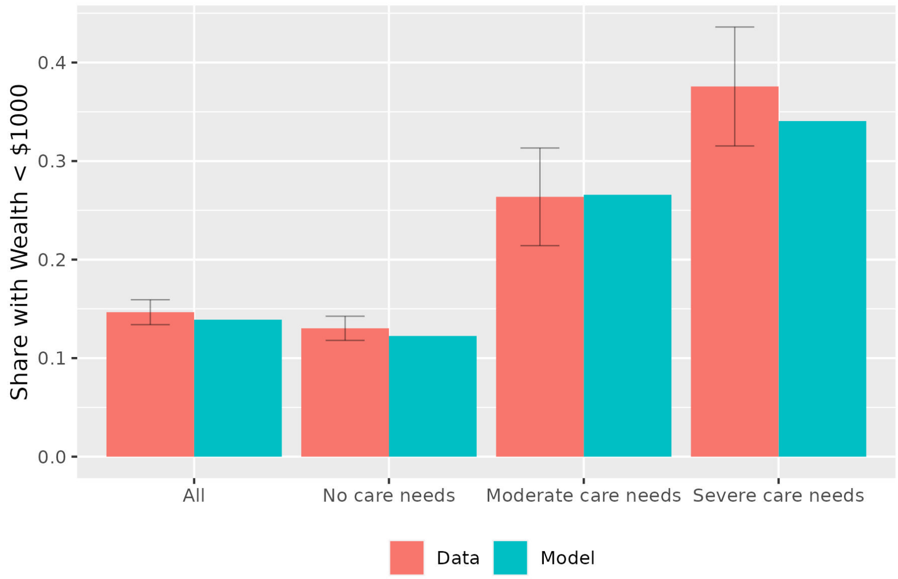

Notes: standard errors clustered at household level. For the four cases, N in the real data is 16888, 15414, 766 and 708 respectively.

The figure suggests that the model is able to match patterns of drawdown of wealth as parents develop care needs.

In Appendix 1.H, I present further figures showing that the model is also able to broadly match levels of children living near parents by gender and parental health status.

1.5 Counterfactuals

In this section I use the estimated model to evaluate some counterfactuals of interest. I focus on two cases: analysing the relative contributions of wages and preferences to the gender care gap, and quantifying the importance of strategic interaction in exacerbating the gender care gap. The first of these exercises gives insight into how much of a narrowing of the gender care gap we can expect from a narrowing of the gender wage gap. The second of these exercises is useful because of what it can tell us about the likely consequences of increasing numbers of one-child families - for whom strategic interaction between children is irrelevant - on the gender care gap, as well as the extent to which changes in one family member's incentives to provide care will affect the care provision of other family members.

1.5.1 Wages vs. preferences in the gender care gap

In the estimated model, daughters differ from sons in two fundamental ways. First, daughters are less averse to providing care[39], and daughters face lower opportunity costs of providing care as their wages are lower. Both of these drive daughters to provide more care than sons in the model, as in the data. It is interesting to consider which of these is the stronger driver of the gender care gap.

To do this, I consider sons' and daughters' care rates in the counterfactual case where there are no preference differences between sons and daughters. In particular, I set sons preferences for providing care so that they are the same as daughters' preferences in the original model[40].

Evaluating counterfactuals of this nature is difficult in a model with multiple equilibria. To make any progress, I must make some assumptions about the counterfactual equilibrium that is selected when parameters change. I adapt the approach of Aguirregabiria and Ho (2012) in imposing smoothness assumptions on the equilibrium selection mechanism and using a Taylor approximation around the original estimated equilibrium to approximate the new counterfactual equilibrium when policies change. More detail on these steps is given in Appendix 1.G.

The results of the counterfactual analysis are shown in Figure 1.5. The figure shows the care rates of sons and daughters, as well as the percentage gender care gap[41], both in the baseline model (left-hand set of bars) and the counterfactual of interest (right-hand set of bars).

Figure 1.5 shows that if sons had the same preferences over care as daughters, then their care rate would rise significantly, and daughters' care rate would drop slightly[42]. The net effect of this is that parents receive more care from their children - as sons are less averse to providing care - and the gender care gap shrinks from 56% (i.e. daughters are 56% more likely to provide care than sons, conditional on having a parent with care needs) to 11%, a decrease of 45 percentage points, or 81% of the original gender care gap. In other words, around four-fifths of the gender care gap is explained by unobserved preference differences. In Appendix 1.H I back up this finding by considering an alternative case where preferences are fixed but the gender wage gap is equalised. In that case, the change in the gender care gap is negligible[43].

[39] Though, as discussed in Section 1.3 above, the parameter on daughter's utility net benefit of providing care captures not only the daughter's attitude to the provision of care but also factors like social norms, so daughters might be just as averse as sons to providing care, but face greater social pressure to do so.

[40] Specifically, I set $\gamma_{son} = \gamma_{son \times younger} = \gamma_{dad \times son} = 0$.

[41] Namely, how much more likely daughters are to provide care relative to sons, expressed in percentage terms.

[42] The daughters' care rate drops because, if they have brothers, they respond strategically to their brothers providing more care by providing less care themselves.

[43] Note that even if there were no preference differences and no wage differences between sons and daughters in the model there would likely still be some gender care gap because of other covariates

Figure 1.5: Gender care gap - no preference difference between sons and daughters

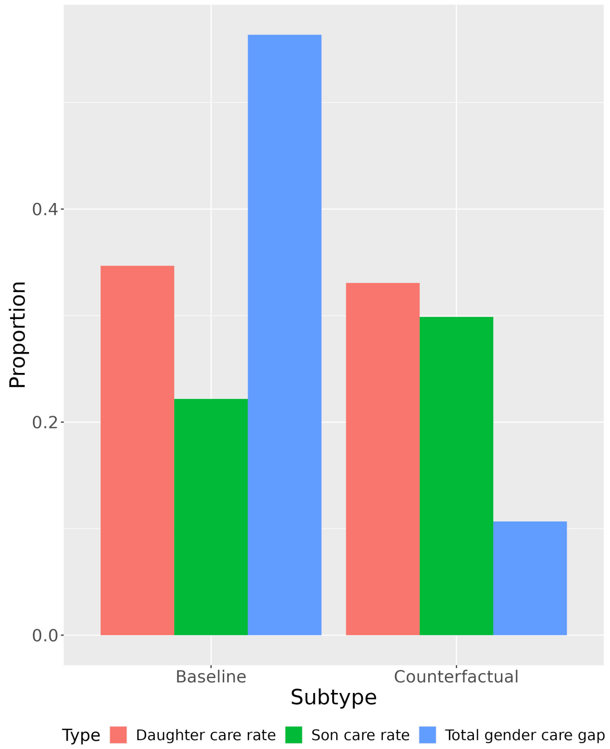

Notes: "care rate" is the probability of providing care conditional on having parent with care needs. "Total gender care gap" how much larger, in percentage terms, daughters' care rate is relative to sons' care rate.

This is a striking and somewhat surprising result. Mechanically, it is driven by the low preference for leisure estimated in the model. The wage gap is relevant to the care decision only to the extent that work choices and care choices correlate with each other: if work and care choices were made entirely independently of one another, the wage would have no impact on a child's incentive to provide care. Moreover, work choices and care choices correlate with each other only to the extent that a child values leisure, and both working and providing care eat into their leisure time. If there were no time cost associated with providing care, or more broadly if the child suffered no loss of leisure utility from working and caring at the same time, then the care decision would be independent of the work decision. Thus, the fact that the estimate for preference for leisure is low dampens the importance of the gender wage gap in driving the gender care gap.

This being said, there is some prima facie evidence in the data to suggest that opportunity cost differentials between children in the same family do not drive differences in care provision. The fixed-effects regressions in Table 1.1 in Section 1.2 showed that using parent-wave fixed effects, a child with more education - hence, one might assume, a higher opportunity cost of providing care in terms of foregone wages - relative to their

which correlate both with gender and with the probability of providing care, such as living close to one's parents, or not having worked in the previous period at the start of the model.

siblings is in fact more likely to provide care. In this light, the fact that gender wage gaps explain so little of the gender care gap is less surprising[44].

Two remarks are in order on this counterfactual. First, the fact that so much of the gender care gap is driven by unobservable preference differences is in one sense discouraging because we have little insight into the nature and causes of this preference difference. However, in a different sense it is important and interesting that so little of the gender care gap is driven by differences in opportunity costs as this suggests that even if the gender wage gap shrinks, there will not necessarily be a significant reduction in the gender care gap, as the causes of the gender care gap are more deep-rooted.

Second, just because the gender care gap is largely driven by differences in preferences for providing care does not mean it is necessarily benign, in the sense that plausibly it is benign for people who have a taste for doing a task to self-select into performing that task. The preference differences here capture not only daughters' enjoyment, or lack of burden, of providing care, but also the extent to which daughters might suffer from going against parental preferences or societal gender norms. These latter two forms of preference difference are plausibly less benign. Thus, just because the gender care gap is largely driven by preference differences does not mean that it is not a matter of policy interest.

1.5.2 The role of strategic interaction

I also examine how big a role strategic interaction plays in the gender care gap. If sons' care is crowded out by their sisters' care, it is interesting to consider how much smaller the gender care gap would be - i.e. how much more care sons would provide, relative to daughters - if the sons had no sisters to rely on to "step up" and provide the care.

To do this, I use the model to generate simulated choices while imposing that child i's sibling never provides any care, and that child i knows this to be the case. This counterfactual captures what might happen to child i's caring behaviour if child i's sibling made a binding commitment to provide no care, or if they moved away to the other side of the world, or if they died. In other words, each child in this case is acting as if they are an only child at every point that they make a decision, as they ignore the possibility that their sibling could provide any care.

The results are presented in Figure 1.6 below. The left-hand panel assesses the change in the gender care gap for the whole population. The right-hand panel assesses the change in the gender care gap only for those families where in the original data both children start the model living close to their parents.

For the whole population, both female and male care rates (i.e. the probability of

[44] Groneck (2017) similarly notes that there is little evidence that children with higher opportunity costs are less likely to provide care.

Figure 1.6: Gender care gap with no interaction

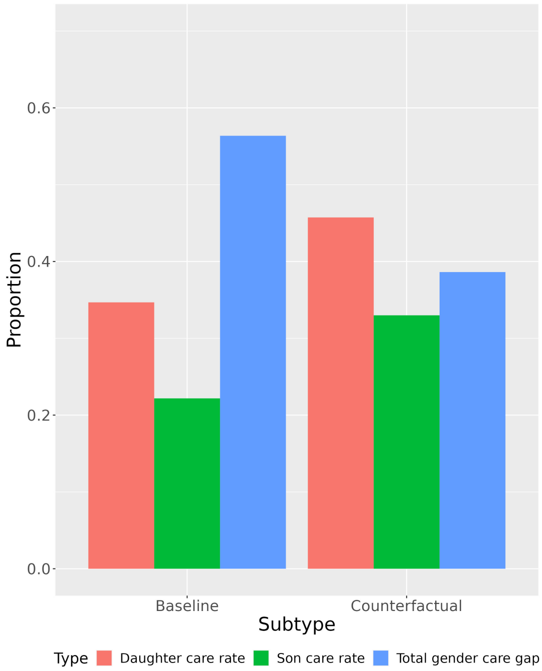

(a) All families

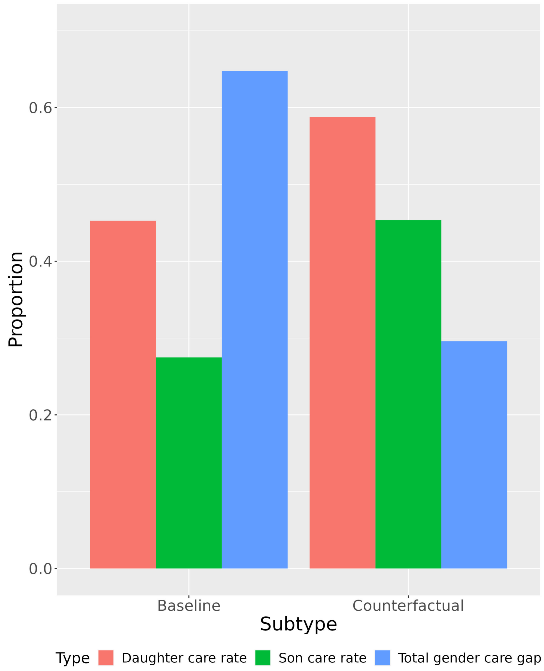

(b) Families where both children start close

Notes: "care rate" is the probability of providing care conditional on having a parent with care needs. "Total gender care gap" how much larger, in percentage terms, daughters' care rate is relative to sons' care rate.

providing care given one's parent has care needs) increase relative to baseline: all children provide more care because they no longer have a sibling to "step up" and provide care if they shirk, with the result that parents receive more care from their children. The proportional increase is greater for men than for women, so the gender care gap shrinks by around 18 percentage points, or 31%.

For the population of families where both children live near their parent at the start of the model, the effect is stronger: the gender care gap shrinks by 35 percentage points, or around 54%. The effect of strategic interaction is stronger in these cases because the existence of a sibling is more relevant to child i's decisionmaking. If both child i and child i's sibling live close to a parent, then it is more likely that child i's sibling would provide care if child i shirks. Hence, the sudden absence of child i's sibling in the counterfactual makes a bigger difference to child i's decisionmaking than if child i's sibling lived far away.

Note that the effect of different preferences by gender on the gender care gap overlaps with the effect of strategic interaction on the gender care gap. This is because in the no-different-preferences counterfactual, I take account of agents' strategic responses to the change in parameters, which will amplify the initial effect of the change in parameters[45].

[45] Note, for example, that in Figure 1.5 the female care rate decreases slightly. This is because daughters

1.6 Conclusion

In this paper I presented a dynamic discrete-choice model of strategic interaction between siblings in providing care to a common elderly parent. The model allowed for endogenous location and work status by letting each child make care, work and location choices every period, and the model isolated the role of interdependence of children's care decisions by separately estimating the utility each child receives from providing care themselves and the utility they receive from having their parent cared for by anyone. By considering non-cooperative interaction between sets of siblings over time, and allowing multiple dimensions of choice for each child in each period, the model goes beyond existing models in the literature and allows a detailed examination of what drives differences in care provision within families, particularly by gender of the child.

The results of the estimation suggest that having a parent receive care is a public good for all children yet it is costly for each child to provide care. Hence, children strategically shirk and underprovide care relative to how much they would provide if they did not have a sibling. Sons shirk relatively more than daughters and strategic interaction of this nature explains around 31% of the total gender care gap. A policy implication of this result is that policies which change agents' costs and preferences for providing care could have significant effects on the overall level and distribution of care given that other agents will respond strategically and adjust their own care levels.

Also, the results of the estimation show that it is unobserved preference differences for providing care between sons and daughters, rather than differences in opportunity costs through the gender wage gap, that are the chief driver of differences in care rates between sons and daughters. Thus, policies which successfully close the gender wage gap in isolation would not substantially change the division of caring responsibilities between sons and daughters.

For reasons of tractability I have significantly simplified the decision set of children in the model (e.g. they do not save, and only have binary location, work and care choices), and have omitted the parent as a player altogether. Future work might consider including the parent - or, ideally, the parent and their spouse, if any - and allowing all players a richer choice set, to see if the conclusions of the more restricted model presented here still hold.

Moreover, the absence of a model of parental preferences or the effect of caregiving on parental outcomes means that this paper is not equipped to offer an in-depth analysis of the efficiency of care arrangements, because doing so would require an understanding of the benefits of different kinds of care arrangements above and beyond their impact on children. Further development of the ideas of this paper along this direction would

respond to sons raising their care effort by cutting their own care effort, which would not be the case if I were ignoring strategic interaction.

allow a more considered analysis of drivers of inefficiency in this area and potential policy interventions to alleviate this inefficiency.

Finally, this paper leaves unresolved the exact drivers of preference differences between sons and daughters when it comes to the provision of care. Future work could decompose this preference difference, into (for instance) the enjoyment, or lack of burden, daughters derive from providing care versus the burden daughters experience from going against parental preferences or social norms, while maintaining the setting of dynamic interaction. This would allow a more sophisticated treatment of what policy approaches, if any, are to be used to reduce the gender care gap.

References

Aguirregabiria, Victor, Allan Collard-Wexler, and Stephen P. Ryan (2021). "Dynamic Games in Empirical Industrial Organization". In: *Handbook of Industrial Organization*. arXiv. DOI: 10.48550/arXiv.2109.01725.

Aguirregabiria, Victor and Chun-Yu Ho (2012). "A dynamic oligopoly game of the US airline industry: Estimation and policy experiments". In: *Journal of Econometrics*. The Econometrics of Auctions and Games 168.1, pp. 156–173. ISSN: 0304-4076. DOI: 10.1016/j.jeconom.2011.09.013.

Aguirregabiria, Victor and Pedro Mira (2007). "Sequential Estimation of Dynamic Discrete Games". In: *Econometrica* 75.1, pp. 1–53. ISSN: 1468-0262. DOI: 10.1111/j.1468-0262.2007.00731.x.

— (2019). "Identification of games of incomplete information with multiple equilibria and unobserved heterogeneity". In: *Quantitative Economics* 10.4, pp. 1659–1701. ISSN: 1759-7331. DOI: 10.3982/QE666.

Angrist, Joshua D. and William N. Evans (1998). "Children and Their Parents' Labor Supply: Evidence from Exogenous Variation in Family Size". In: *The American Economic Review* 88.3. Publisher: American Economic Association, pp. 450–477. ISSN: 0002-8282.

Bajari, Patrick, C. Lanier Benkard, and Jonathan Levin (2007). "Estimating Dynamic Models of Imperfect Competition". In: *Econometrica* 75.5, pp. 1331–1370. ISSN: 1468-0262. DOI: 10.1111/j.1468-0262.2007.00796.x.

Barczyk, Daniel and Matthias Kredler (2018). "Evaluating Long-Term-Care Policy Options, Taking the Family Seriously*". In: *The Review of Economic Studies* 85.2, pp. 766–809. ISSN: 0034-6527. DOI: 10.1093/restud/rdx036.

Bom, Judith et al. (2019). "The Impact of Informal Caregiving for Older Adults on the Health of Various Types of Caregivers: A Systematic Review". In: *The Gerontologist* 59.5, e629–e642. ISSN: 1758-5341. DOI: 10.1093/geront/gny137.

Byrne, David et al. (2009). "Formal home health care, informal care, and family decision making". In: *International economic review (Philadelphia)* 50.4. Place: Malden, USA Publisher: Blackwell Publishing Inc, pp. 1205–1242. ISSN: 0020-6598. DOI: 10.1111/j.1468-2354.2009.00566.x.

Checkovich, Tennille J. and Steven Stern (2002). "Shared Caregiving Responsibilities of Adult Siblings with Elderly Parents". In: *The Journal of Human Resources* 37.3, p. 441. ISSN: 0022166X. DOI: 10.2307/3069678.

Cleveland Clinic (2024). *Activities of Daily Living*. URL: https://my.clevelandclinic.org/health/articles/activities-of-daily-living-adls.

Fontaine, Roméo, Agnès Gramain, and Jérôme Wittwer (2009). "Providing care for an elderly parent: interactions among siblings?" In: *Health Economics* 18.9, pp. 1011–1029. ISSN: 1099-1050. DOI: 10.1002/hec.1533.

Genworth (2024). *Cost of Long Term Care by State — Cost of Care Report — Genworth*. URL: https://www.genworth.com/aging-and-you/finances/cost-of-care.html (visited on 11/20/2024).

Grigoryeva, Angelina (2017). "Own Gender, Sibling's Gender, Parent's Gender: The Division of Elderly Parent Care among Adult Children". In: *American Sociological Review* 82.1. Publisher: SAGE Publications Inc, pp. 116–146. ISSN: 0003-1224. DOI: 10.1177/0003122416686521.

Groneck, Max (2017). "Bequests and Informal Long-Term Care: Evidence from HRS Exit Interviews". In: *Journal of Human Resources* 52.2. Publisher: University of Wisconsin Press, pp. 531–572.

HRS (2024). *Health and Retirement Study, public use dataset.* Products used: RAND HRS Longitudinal File 2020 (V1), RAND HRS Longitudinal File Version M, RAND HRS Longitudinal File Version N, RAND HRS Family Data 2018 (V1), HRS Exit and Post-Exit Files 2002-2018, RAND HRS Fat Files 1998-2018. Produced and distributed by the University of Michigan with funding from the National Institute on Aging (grant numbers NIA U01AG009740 and NIA R01AG073289). Ann Arbor, MI.

Hiedemann, Bridget, Michelle Sovinsky, and Steven Stern (2018). "Will You Still Want Me Tomorrow? The Dynamics of Families' Long-Term Care Arrangements". In: *Journal of Human Resources* 53.3. Publisher: University of Wisconsin Press, pp. 663–716. ISSN: 0022166X. DOI: 10.3368/jhr.53.3.0213-5454R1.

Ko, Ami (2022). "An Equilibrium Analysis of the Long-Term Care Insurance Market". In: *The Review of Economic Studies* 89.4, pp. 1993–2025. ISSN: 0034-6527. DOI: 10.1093/restud/rdab075.

Konrad, Kai A. et al. (2002). "Geography of the Family". In: *American Economic Review* 92.4, pp. 981–998. ISSN: 0002-8282. DOI: 10.1257/00028280260344551.

Maruyama, Shiko and Meliyanni Johar (2017). "Do siblings free-ride in "being there" for parents?" In: *Quantitative Economics* 8.1, pp. 277–316. ISSN: 1759-7331. DOI: 10.3982/QE389.

Mommaerts, Corina (2024). "Long-Term Care Insurance and the Family". In: *Journal of Political Economy*. Publisher: The University of Chicago Press, pp. 000–000. ISSN: 0022-3808. DOI: 10.1086/732887.

PSID (2024). *Panel Study of Income Dynamics, public use dataset*. Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan. Ann Arbor, MI.

Rainer, Helmut and Thomas Siedler (2009). "O Brother, Where Art Thou? The Effects of Having a Sibling on Geographic Mobility and Labour Market Outcomes". In: *Economica* 76.303. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0335.2008.00696.x, pp. 528–556. ISSN: 1468-0335. DOI: 10.1111/j.1468-0335.2008.00696.x.

RAND (2024). *RAND HRS*. Products used: RAND HRS Longitudinal File 2020 (V1), RAND HRS Longitudinal File Version M, RAND HRS Longitudinal File Version N, RAND HRS Family Data 2018 (V1), RAND HRS Fat Files 1998-2018. Produced by the RAND Center for the Study of Aging, with funding from the National Institute on Aging and the Social Security Administration. Santa Monica, CA.

Skira, Meghan M. (2015). "Dynamic Wage and Employment Effects of Elder Parent Care". In: *International Economic Review* 56.1. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University], pp. 63–93. ISSN: 0020-6598.

Sonnega, Amanda and David Weir (2014). "The Health and Retirement Study: A Public Data Resource for Research on Aging". In: *Open Health Data* 2. DOI: 10.5334/ohd.am.

Sovinsky, Michelle and Steven Stern (2016). "Dynamic modelling of long-term care decisions". In: *Review of Economics of the Household* 14.2. Number: 2 Publisher: Springer New York LLC, pp. 463–488. ISSN: 1569-5239. DOI: 10.1007/s11150-013-9236-3.

Stern, Steven (2023). "Where Have All My Siblings Gone?" In: *Journal of Human Resources* 58.3. Publisher: University of Wisconsin Press Section: Articles, pp. 852–892. ISSN: 0022-166X, 1548-8004. DOI: 10.3368/jhr.59.1.0220-10739R2.

Van Houtven, Courtney Harold, Norma B. Coe, and Meghan M. Skira (2013). "The effect of informal care on work and wages". In: *Journal of Health Economics* 32.1, pp. 240–252. ISSN: 0167-6296. DOI: 10.1016/j.jhealeco.2012.10.006.

Weir, David, Jessica Faul, and Kenneth Langa (2011). "Proxy interviews and bias in the distribution of cognitive abilities due to non-response in longitudinal studies: a comparison of HRS and ELSA". In: *Longit. Life Course Stud.*

Wettstein, Gal and Alice Zulkarnain (2017). *How much long-term care do adult children provide*. Tech. rep. Center for Retirement Research at Boston College.

Appendices

1.A Supplementary descriptive figures and tables

1.A.1 HRS sample descriptive statistics

Table 1.6 presents some key descriptive statistics on the full sample of HRS respondents for Waves 4 to 14. This is the sample used to establish the stylised facts in Section 1.2. I also include statistics for the subsample of single respondents and the subsample of single respondents with exactly two children because the estimation will focus on these groups.

Table 1.6: Descriptive statistics

	All	Single respondents	Single respondents, 2 kids
Female	0.59	0.76	0.75
Age	67.08	71.27	70.77
Couple	0.67	0.00	0.00
# Kids	3.38	3.18	2.00
- of which live within 10 miles	0.97	1.00	0.65
White Caucasian	0.76	0.69	0.74
Homeowner	0.76	0.57	0.61
Some college	0.43	0.36	0.42
Wealth ($1000s)	117	73	107
Observations	200 385	67 123	17 812

Notes: all statistics are weighted means using HRS respondent-level weights, apart from wealth, which is a weighted median to reduce the impact of outliers. As wealth is measured at the household rather than individual level, I allocate half of each couple household's wealth to each member of the couple.

The overall HRS sample skews female, given lower life expectancy for men. About two-thirds of respondents are in a couple, about three-quarters are homeowners and about one-half had at least some college education.

The subsample of singles is notably more female, slightly older, and less likely to own a home or to have gone to college. Median wealth is considerably less than in the full sample. Single respondents are also slightly less likely to be white. It is notable that the average number of kids is lower for single respondents, yet the average number of kids living within 10 miles is higher, suggesting that a higher proportion of the children of single parents live near their parent.

Finally, the subsample of single respondents with exactly two kids is slightly richer and more educated than the subsample of all single respondents. Indeed, along the wealth, education and race dimensions this subsample more closely resembles the full HRS sample than it does the subsample of singles.

1.A.2 Hours of care received by care source

Figure 1.7 presents information on the sources of care to elderly people. Each bar shows mean hours of care received by a given parent with care needs by each of five possible sources: spouses, sons, daughters, other family members and non-family sources. I calculate this separately for the whole sample, the subsample of single respondents and the subsample of single respondents with only two children. Note that I do not condition on positive care provision here so the graph captures both the intensive and extensive margins of care provision. Also, employees of institutions are excluded from calculations of total care receipt in the HRS, so the "Non-family" care does not include e.g. care provided by workers in a nursing home.

Figure 1.7: Hours of care per month by care source

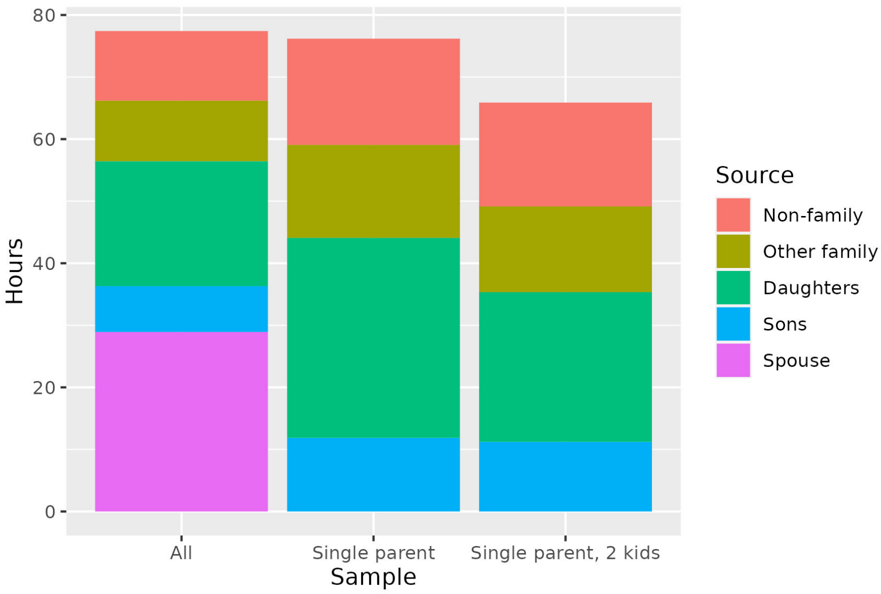

Notes: parents with care needs (ADLs or IADLs >0) only. "Non-family" care excludes care from employees of institutions. Means weighted by HRS respondent-level weights.

For the sample as a whole, spouses provided the biggest share of care. A given parent with care needs will receive around 29 hours of care per month from a spouse, compared to 7 hours from any sons they have, 20 hours from any daughters they have, 10 hours from other family members and 11 hours from non-family sources, adding up to a total of around 77 hours of care.

For the sample of single respondents only, it is notable that the average amount of care hours received stays approximately the same. In other words, all of the other sources of care increase their output to compensate for the lost care from the spouse. In particular, daughters are the biggest single source of care, and provide much more care than sons.

Comparing this to the sample of single respondents with two kids, the sum of care hours from children decreases somewhat, leading to a decrease in the total amount of care received, simply because there are fewer candidate children to provide the care. It is notable that the difference between daughters' care hours and sons' care hours is less pronounced when there are only two children. This is because in many two-child families there are no daughters, hence sons cannot leave it to their sisters to provide the care so must "step up" themselves.

Note that despite the appearance of the graph, the typical elderly individual with care needs does not receive care from many sources in roughly equal amounts. Instead, most elderly people will have at most one primary caregiver. In the full sample, 75% of people receiving help with IADLs, ADLs or finances from at least one person receive help from exactly one person.

1.A.3 Extra descriptive regressions and graphs

1.A.3.1 Household care hours regression

Table 1.7: Predictors of care provision and the amount of care provided - household level

	Dependent variable:		
	I(HH hours help p.m. > 0)		log(HH Hours help p.m)
	(1)	(2)	(3)
Kid is daughter	0.017***	0.020***	0.158**
	(0.003)	(0.004)	(0.055)
Parent is mother	0.012***		0.019
	(0.003)		(0.051)
Daughter x Mother	0.051***	0.061***	0.164**
	(0.003)	(0.004)	(0.064)
Kid is eldest	0.007***	0.007	0.005
	(0.002)	(0.004)	(0.034)
Kid is youngest	0.020***	0.007	−0.017
	(0.002)	(0.004)	(0.033)
Kid lives ≤ 10 miles away	0.154***	0.180***	0.559***
	(0.002)	(0.003)	(0.030)
Kid has sister	−0.025***		−0.110**
	(0.002)		(0.034)
Parent in couple	−0.065***		−0.163***
	(0.002)		(0.032)
Kid in couple	−0.011***	−0.013***	−0.285***
	(0.002)	(0.003)	(0.030)
Kid went to college	0.018***	0.027***	−0.192***
	(0.002)	(0.003)	(0.028)
Parent × wave FEs	N	Y	N
Observations	117 306	117 333	13 769
Adjusted R^2	0.178	0.309	0.113
Mean dep. var.	0.099	0.099	3.043

Notes: estimation via OLS. Dependent variable is household care hours by the kid, i.e. the sum of kid and their partner's care hours. Only children of parents with care needs (ADL or IADL >0) are included. Regression weights are HRS respondent-level weights. In all regressions, other controls are dummies for number of ADL and IADL difficulties, education of parent, number of total kids of the parent, total number of grandkids of the kid, quadratics in age for kid and parent, a dummy for parent being White Caucasian, log of parent wealth and homeownership status of kid. For the fixed effects regression, the parent-level controls drop out. Standard errors clustered at the household level. *p<0.05; **p<0.01; ***p<0.001.

1.A.3.2 Logit regression

Table 1.8: Average marginal effects from logit models

	Dependent variable:
	Logit AME
	(1)
Kid is daughter	0.060***
	(0.003)
Parent is mother	0.042***
	(0.003)
Daughter x Mother	0.047***
	(0.005)
Kid is eldest	0.003
	(0.003)
Kid is youngest	0.008*
	(0.003)
Kid lives $\leq$ 10 miles away	0.143***
	(0.003)
Kid has sister	−0.014***
	(0.004)
Parent in couple	−0.054***
	(0.003)
Kid in couple	−0.025***
	(0.003)
Kid went to college	0.007*
	(0.003)
Parent $\times$ wave FEs	N
Pseudo R^2	0.255
Observations	128 878
Mean dep. var.	0.094

Notes: estimation via logit regression. Quantities in the table are average marginal effects. Only children of parents with care needs (ADL or IADL >0) are included. Regression weights are HRS respondent-level weights. In all regressions, other controls are dummies for number of ADL and IADL difficulties, education of parent, number of total kids of the parent, total number of grandkids of the kid, quadratics in age for kid and parent, a dummy for parent being White Caucasian, log of parent wealth and homeownership status of kid. Standard errors clustered at the household level, constructed using the delta method. *p<0.05; **p<0.01; ***p<0.001.

1.A.3.3 Gender care gaps controlling for location

Figure 1.8: Probability of providing care by type of kid - controlling for location

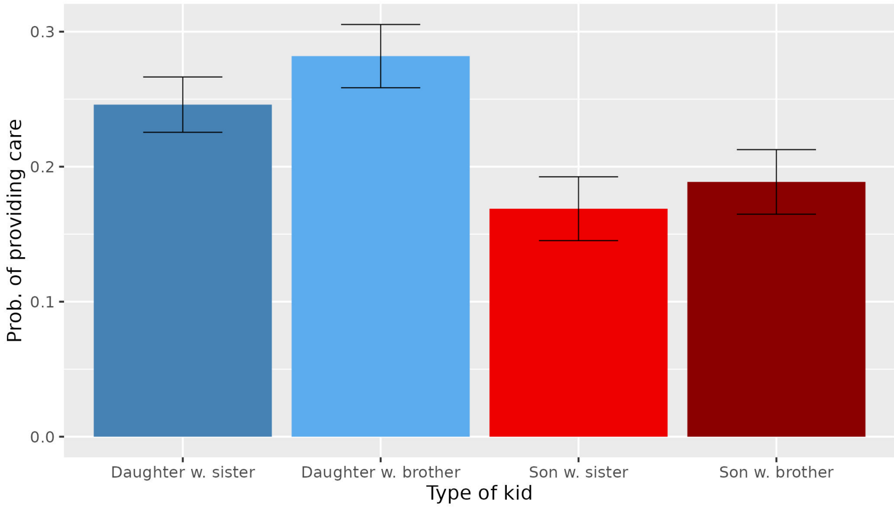

Notes: probabilities of providing care to a single parent with care needs for daughters whose only sibling is a sister, daughters whose only sibling is a brother, sons whose only sibling is a sister and sons whose only sibling is a brother, controlling for location (whether within 10 miles of the parent) of both siblings through an OLS regression with location and family type as RHS controls. Probabilities presented are the regression coefficients on family type dummies, added to an intercept calculated as the sum of the coefficients on the location dummies multiplied by mean propensity to live within 10 miles of the parent, for both children respectively. Regression weights are HRS respondent-level weights.

1.A.3.4 Gender care gaps across family type

Table 1.9 presents measures of care provision by sons and daughters in families of different types.

Table 1.9: Ratios of daughters' care to sons' care by family type

	I(Provides care)			Care hours		
	Son	Daughter	Ratio	Son	Daughter	Ratio
One-child family	0.17	0.27	1.57	12.85	26.88	2.09
Two-child daughter+son family	0.10	0.19	1.91	5.90	15.24	2.58
Two-child D+D family	-	0.17	-	-	13.26	-
Two-child S+S family	0.12	-	-	6.20	-	-

Notes: only parents with care needs (difficulties with at least one ADL or IADL) included. "I(Provides care)" is a dummy for whether the child provides any care, "Care hours" is a measure of care hours per month. "Ratio" is daughters' mean outcome variable divided by sons' mean outcome variable - i.e. a measure of how much more care daughters provide than sons. Means weighted by HRS respondent-level weights.

1.A.3.5 Caring and in-laws

So far I have been ignoring the fact that a married couple has in general two sets of parents that they could care for: in heterosexual couples, the husband's parents and the wife's parents. Moreover, if those two sets of parents live far from each other, each married couple can choose only one to care for. How exactly a married couple addresses this problem has significant welfare implications. For instance, it could be that a married couple decides which set of parents to care for depending on the probability of each set of parents receiving care from other sources, in which case there is some redistribution of care effort to those who need it most. A full examination of this complicated issue is beyond the scope of this paper and in any case would require more in-depth data than the HRS provides on siblings-in-law and parents-in-law. However, some preliminary steps can be taken in this direction.

It has already been established above that a given child X is more likely to provide care to their parent if they have no sisters to "step up" in their absence. However, to examine how married couples make decisions of who to care for, we need to consider whether child X is less likely to provide care to their parent if their spouse Y does not have any sisters. If this were the case, it would suggest that X and Y are deciding which set of parents to care for depending on how likely they are to receive care from other sources.

Table 1.10 below assesses the importance of this channel. It reports results of a linear probability model with the LHS variable being a dummy for whether an HRS respondent[46] ("R") lives within 10 miles of their mother, and the key RHS variables being dummies for whether R and their spouse respectively have any sisters and the total number of siblings that R and their spouse have[47]. I restrict the sample to only those married HRS respondents with living mothers and whose spouses have at least one parent alive. Column 1 shows results for the whole sample and Column 2 limits the regression to only male HRS respondents.

People are less likely to live within 10 miles of their mother (and thus be in a position to provide care) if they have a sister, and if they have many siblings. However, the association between R's location decision and the family structure of R's spouse is less clear cut. It seems that R is more likely to live near their mother if R's spouse has many siblings - which is consistent with R and R's spouse deciding to live near R's mother rather than the spouse's parent if they believe the spouse's parent will be well cared for anyway - but the association is weak.

[46]Note that this regression uses data on HRS respondents' location relative to their parents, rather than the location of the children of HRS respondents relative to HRS respondents themselves, as is the setting for the rest of the paper. This is because there is no information in the survey on the families of the spouses of the children of HRS respondents.

[47]Other RHS regressors are R's gender (in the regression with both male and female respondents), homeownership status, whether white Caucasian, education and a polynomial in age.

Table 1.10: Caring and in-laws

	Dependent variable:	
	R lives within 10 miles of R's mother	
	(1)	(2)
R has sister	−0.034**	−0.057***
	(0.011)	(0.016)
R's total siblings	−0.006**	−0.003
	(0.002)	(0.003)
R's spouse has sister	−0.002	0.001
	(0.011)	(0.015)
R's spouse total siblings	0.004*	0.006*
	(0.002)	(0.003)
Observations	12 791	6464
Adjusted R^2	0.013	0.013
Mean dep. var.	0.347	0.352

Notes: estimation via OLS. "R" stands for HRS Respondent. Column 1 is for all Respondents, Column 2 is for men only. Other controls are R's gender (Column 1 only), homeownership status, whether white Caucasian, education and a polynomial in age. Regression weights are HRS respondent-level weights. *p<0.05; **p<0.01; ***p<0.001.

As such, while the issue of links between the provision of care in different nuclear families through in-law relationships is important and deserving of further attention, it does not appear to be a first-order issue for the purposes of this paper, so I will largely leave it aside.

1.A.3.6 Care received by family composition

Figure 1.9 shows the mean child-provided care hours received by sick single parents of two children depending on the gender composition of their children. Parents with more daughters clearly receive higher overall child-provided care hours than parents with fewer daughters.

Figure 1.9: Mean care hours received by composition of kids

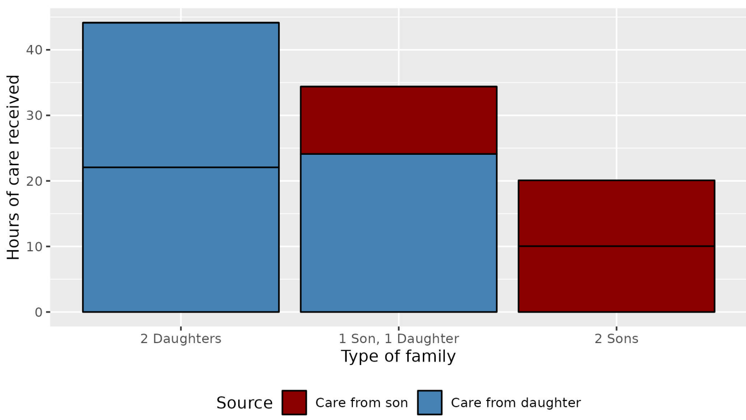

Notes: parents with care needs (ADLs or IADLs >0) only. Means weighted by HRS respondent-level weights.

1.B Children's location and parental health

To examine the effect of health shocks to parents on kids' location decisions I carry out an event study using the respondent-kid data in the HRS. I run the following regression:

$$Y_{ijt} = \alpha_{ij} + \kappa_t + \sum_{k \in -3,...2} \beta_k I(Period_{jt} = k) + X_{ijt}\gamma + \epsilon_{ijt} \tag{1.20}$$

In this regression, Y_{ijt} is the outcome variable, $Period_{jt}$ captures the period relative to the period where the parent j develops care needs for the first time (i.e. in that period, $Period_{jt} = 0$). X_{it} is a set of time-varying controls[48]. I include child $\times$ parent fixed effects α_{ij} and wave fixed effects κ_t.

[48]In this case, I control for polynomials in age of parent and child and couple status of parent and child.

Figure 1.10 plots the results of this estimation. I estimate Equation 1.20 separately for sons and daughters. The height of each line at period k is the value of β_k estimated from the regressions.

In the left-hand panel the outcome variable is a dummy for whether child i lives within 10 miles of a parent j at t. In the right-hand panel, to avoid concerns about it being parents who choose which child to move close to, the outcome variable is a dummy for whether child i moves back to live with their parent in t. Note that in my estimation sample in only 19% of cases where a parent lives more than 10 miles from a given child in $t-1$ and lives within 10 miles in t, it is the parent who has moved, not the child. In other words, in roughly four fifths of cases it is the children who move close to the parent rather than the other way round[49].

Figure 1.10: Children's location before and after health shocks

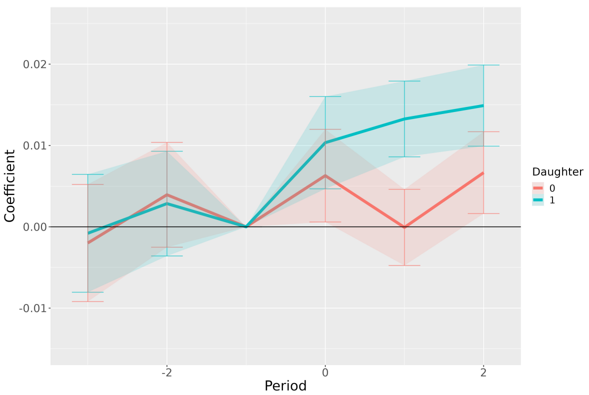

(a) Outcome var: kid lives within 10 miles of parent

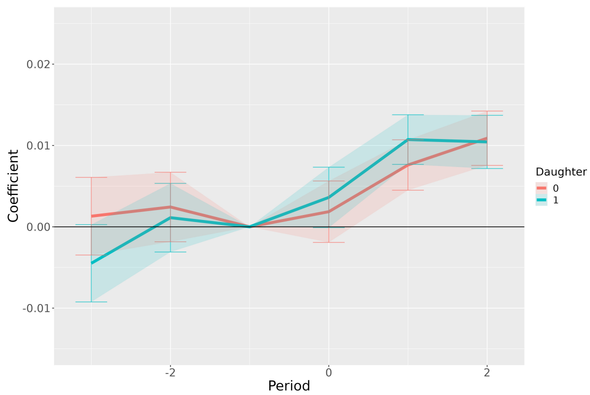

(b) Outcome var: kid moves to live within 10 miles of parent

Notes: estimation via OLS. Standard errors clustered at child × parent level. Confidence bands show 95% confidence intervals. Coefficient β_{-1} normalised to 0.

The figures suggest that children are more likely to live closer to their parent after a parent falls ill, with this estimated association being larger for daughters. To put the magnitudes of the coefficients into context, in the estimation sample the proportion of kids who live within 10 miles of the relevant parent was 37.6% for sons and 38.7% for daughters.

Thus, there is some evidence to suggest that children do move back to live near their parents when their parents develop care needs. This is particularly true of daughters.

[49] If we consider only cases where the parent and child live within 10 miles of each other in t and not $t-1$, *and* parent has developed care needs for the first time in the data in t, then in 21% of these cases it is the parent who moves.

1.C Other possible mechanisms of compensation

1.C.1 Extensive margin of bequest provision

Table 1.11: Bequest receipt by care provision (extensive margin)

	Dependent variable:				
	Whether any bequest received				
	(1)	(2)	(3)	(4)	(5)
$CARE_i$	0.084***	0.049***	0.049***	0.050***	0.051***
	(0.012)	(0.007)	(0.007)	(0.007)	(0.007)
Parent FEs	Y	Y	Y	Y	Y
Observations	6681	12 428	12 303	11 805	11 185
Adjusted R^2	0.671	0.839	0.838	0.830	0.819
Mean dep. Var	0.692	0.372	0.366	0.342	0.313

Notes: estimation via OLS. Dependent variable is a dummy for whether any bequest is received by the child in question. Column 1 considers only those children of parents with a positive estate at death. Column 2 considers all children. Column 3 drops children of parents with estates above the 99th percentile. Column 4 drops children of parents with estates above the 95th percentile. Column 5 drops children of parents above the 90th percentile. Controls are age of child, child education, child income, whether the child owns a home, whether the child lives within 10 miles of the parent, whether the child is co-resident with the parent and frequency of contact with the parent. Regression weights are HRS respondent-level weights. *p<0.05; **p<0.01; ***p<0.001.

1.C.2 Deciles of bequest distribution

Figure 1.11: Deciles of bequest distribution

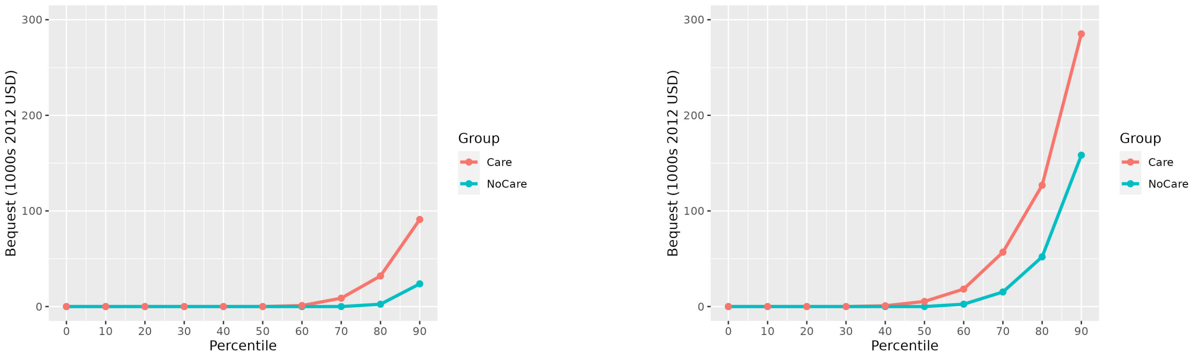

(a) Parent not college-educated

(b) Parent college-educated

Notes: "Care"/"Non-Care" correspond to the deciles of the bequest distribution for those children with $CARE_i = 1$ and 0 respectively. Deciles are weighted by HRS respondent-level weights.

1.C.3 Total estate by care receipt

Table 1.12: Total estate by care receipt

	Total estate (1000s of 2012 USD)
Whether any child provided care	57.290***
	(19.181)
Observations	3193
Adjusted R^2	0.030
Mean dep. var.	193.458

Notes: estimation via OLS. Observations at the parent level. Other RHS controls are dummies for number of children and a quadratic of age of parent at time of death. Regression weights are HRS respondent-level weights. *p<0.05; **p<0.01; ***p<0.001.

1.C.4 Compensation through rent-free accommodation

Some authors (e.g. Barczyk and Kredler (2018)) have argued that in addition to bequests children are compensated through rent-free accommodation when they are providing care to their parents.

HRS respondents are not directly asked whether their co-resident children pay any rent, but can gain more insight by looking a generation upwards: while there are no data in the HRS on whether the children of survey respondents pay rent to the respondents when they live with the respondents, there is information on whether survey respondents themselves pay any rent when they live with their parents.

Table 1.13 presents these results. Column 1 shows proportions of tenure type for all respondents in the sample, where each observation is a respondent-wave combination. Column 2 shows tenure types for the subsample of respondents who provide more than 40 hours of care per month (I call these "major helpers"[50]). Column 3 restricts this further to those major-helping respondents who co-reside with one of their parents.

Table 1.13: Tenure type by amount of care

	All respondents	"Major helpers"	MH+co-reside with parent
Own	0.79	0.81	0.69
Rent	0.17	0.13	0.15
Live rent-free	0.03	0.05	0.14
Other	0.01	0.01	0.02
Observations	172 145	1802	468

Notes: data from HRS Waves 5 to 14 (2000-18). All statistics are weighted means using HRS respondent-level weights.

The results in Column 2 suggest that only 5% of those providing over 40 hours of care per month to a parent live rent-free. It is important to note that the HRS sample will not be fully representative of the sample of all carers so the average HRS respondent will be older, thus likely richer and less in need of rent-free accommodation, than the average carer. However, even with these caveats, the evidence on children receiving rent-free accommodation as compensation for care provision seems mixed.

This is important because if children do not receive rent-free accommodation in exchange for providing care, and do not receive substantively more by way of bequests, then this casts doubt on the suggestion that care hours are provided in exchange for a financial transfer of some kind.

1.C.5 Compensation through childcare

Finally, it might be argued that children provide care to their parents in exchange for childcare (i.e. care of the grandchildren) received from their parents. In other words, there is a dynamic contract between children and parents: in some period, the parent cares for their grandchildren, and in another period, the child provides care to the parent.

To investigate whether this channel of exchange exists I regress a dummy for whether a kid provides care to a parent, conditional on that parent having care needs, on a set of

[50]HRS respondents are asked how many hours of personal care they provided to their parents since the last interview. I take this figure and divide it by 24 to arrive at the amount per month, given interviews are biennial.

explanatory variables, notably including a dummy for whether that kid is ever observed receiving childcare from that parent[51].

Table 1.14 reports the results. In Column 1, I regress the care dummy on the base set of explanatory variables. In Column 2, I add an interaction between the dummy for whether the kid is ever observed receiving childcare from the parent and the dummy for whether the kid lives within 10 miles of the parent. Columns 3 and 4 are the same as Columns 1 and 2 except parent-level fixed effects are used.

Table 1.14: Childcare and caring for parents

	Dependent variable:			
	Kid provides care			
	(1)	(2)	(3)	(4)
Kid ever recd childcare	0.022***	0.039***	0.021***	0.036***
	(0.002)	(0.003)	(0.004)	(0.004)
Kid lives $\leq$ 10 miles away	0.153***	0.164***	0.157***	0.166***
	(0.002)	(0.002)	(0.003)	(0.004)
Kid ever recd childcare $\times$ Kid lives $\leq$ 10 miles away		−0.036***		−0.033***
		(0.004)		(0.007)
Parent FEs	N	N	Y	Y
Observations	101 208	101 208	101 208	101 208
Adjusted R^2	0.174	0.175	0.282	0.289
Mean dep. var.	0.112	0.112	0.112	0.112

Notes: estimation via OLS. In all regressions, other controls are dummies for number of ADL and IADL difficulties, education of kid and parent, number of total kids of the parent, quadratics in age for kid and parent, a dummy for being White Caucasian, log of wealth and a dummy for being eldest/youngest child. For the fixed effects regression, the parent-level controls drop out. Standard errors clustered at the parent level. Regression weights are HRS respondent-level weights. *p<0.05; **p<0.01; ***p<0.001.

The results in Columns 1 and 3 suggest a substantive association between whether a kid ever received help with childcare from a parent and whether the kid provides care to the parent: kids who received childcare are (using the results from Column 3) 2.1pp more likely to provide care, relative to a mean probability of providing care of 11.2%. However, the inclusion of the interaction term in Columns 2 and 4 shows that this association is largely driven by kids who live further than 10 miles from their parent in the current period. For those kids who live further than 10 miles, those who received childcare in the past are (using the results from Column 4) 3.6pp more likely to provide care than those who did not receive childcare. However, for those kids who live within 10 miles of their parents, the difference in probabilities is only $3.6 - 3.3 = 0.3$pp. One explanation

[51]In particular, this dummy variable is equal to one for a given wave if the parent says that they or their spouse spent more than 100 hours caring for their (great-)grandchildren via that child.

for this is that within the group of kids who live more than 10 miles from parents, those who received childcare from parents are more likely to live closer (e.g. 20 miles away) than those who did not. As such, while this basic analysis clearly does not rule out a childcare-eldercare exchange channel, there seems to be little indication that kids provide care to parents in exchange for parents providing childcare for the kids.

1.D Model discussion

1.D.1 $\omega_{warm}(\cdot)$ in full

In full, Equation 1.15 takes the form:

$$\omega_{warm}(\cdot) = \begin{cases} \gamma_{h1} + \gamma_{far} far_t^i + \gamma_{origfar} far_0^i + \gamma_{son} son^i & \text{if } h_t = 1 \ \& \ k_t^i > 0. \\ +\gamma_{dad} dad + \gamma_{start} start_t^i + \gamma_{younger} I(i = B) & \\ +\gamma_{younger \times son} I(i = B) \times son^i & \\ +\gamma_{dad \times son} dad \times son^i, & \\ & \\ \gamma_{h2} + \gamma_{far} far_t^i + \gamma_{origfar} far_0^i + \gamma_{son} son^i & \text{if } h_t = 2 \ \& \ k_t^i > 0. \\ +\gamma_{dad} dad + \gamma_{start} start_t^i + \gamma_{younger} I(i = B) & \\ +\gamma_{younger \times son} I(i = B) \times son^i & \\ +\gamma_{dad \times son} dad \times son^i, & \\ & \\ 0, & \text{otherwise.} \end{cases} \tag{1.21}$$

Here, far_t^i is a dummy for whether child i lives far (i.e. more than 10 miles) from their parent at time t, son^i is a dummy for whether child i is a son, dad is a dummy for whether the parent is a father, $start_t^i$ is a dummy for whether child i is starting providing care this period, and $I(i = B)$ is a dummy for whether child i is the younger of the two siblings.

1.D.2 Unobserved heterogeneity

Any dynamic model of this kind faces an important problem of unobserved heterogeneity (Aguirregabiria, Collard-Wexler, and Ryan 2021). The model imposes that the only form of unobserved heterogeneity takes the form of iid preference shocks independent across agents and across time. In particular, the model does not allow for permanent unobserved heterogeneity, in the form of – for instance – different unobserved levels of affection in the relationship between a particular child and a particular parent.

This could create problems because each child's initial conditions – for instance, their starting location, work and care choice – could be correlated both with this unobserved heterogeneity and with the child's endogenous state variables at the time of making location, work and care choices in period t. For instance, children who feel more affection for their parents may a) be more likely to live near their parent at the start of the game and b) be more likely to provide care when their parent develops care needs in period t. Due to inertia in location choice, a child who starts the model living near their parent is likely to be living near their parent in period t. However, this would lead to the model overestimating the importance of current location in determining the cost of providing care, because the children who currently live near their parents are a selected sample who are more likely to have unobservable characteristics causing them to have a higher preference for providing care.

I address this problem by allowing children to have different "net warm glow" utility from providing care depending on whether they live far from their parent at the start of period 1, i.e. at the very start of the game. This is the role of the $\gamma_{origfar}$ parameter. This is meant to capture the fact that children living near their parent at the start of the game may be systematically different in terms of their relationship with their parent and their preference for providing care to children who live far from their parent. In effect, I am using the child's initial location choice as a proxy for their unobserved type, and I am allowing types to differ in how much "net warm glow" they derive from providing care.

I take a different approach when it comes to previous caring status – in the estimation sample selection, described in more detail in Section 1.4, I drop any observations where the parent has care needs in their first period in the data and where any child provided care in the previous period. It is important to drop these cases because these are families where a child has already selected into caregiving. Using the same example as above, there will be a correlation between unobserved "affection" between parent and child and initial caregiving status, and between initial caregiving status and caregiving at time t. By imposing that every child has the same initial conditions in terms of previous caregiving I am able to model any selection into caregiving as endogenous within the model.

I do not take extra steps to control for agents' starting labour force status both to keep the model tractable through avoiding extra state variables and because it is less clear how initial work status would correlate with preference for providing care.

1.E Parameters calculated outside the model

1.E.1 Care choices, costs and hours

I define child i as providing care in a given period if their parent reports receiving 8 or more hours a month, or 2 hours a week, in help from child i. The reason for setting a non-zero cutoff in care hours to qualify as providing care is that there is significant variation in the number of hours that children are reported as providing, conditional on providing positive hours. As such, to focus on cases where care effort is roughly comparable across children, I count as providing care only those who provide a substantial amount of care, which I set as being 8 or more hours per month[52].

I assume that children who are the only provider of informal care to their parent provide 13 hours per week when their parent has $h_t = 1$ and 25 hours per week when their parent has $h_t = 2$, matching mean hours of care provided by such children in my model estimation sample[53]. I also assume that when two children provide care at the same time, they each provide half of these totals.

The formal care cost values come from Genworth (2024). I assume those parents who have $h_t = 1$ and receive no informal care are forced to pay for a home health aide, either out of pocket or through Medicaid if eligible, who works the same hours over the two years (1352) that a child would provide in informal care, at a cost of \$15.0k per year in 2012 dollars. I assume that parents who have $h_t = 2$ and receive no informal care live in a semi-private room in a nursing home, at a cost of \$76.6k per year.

1.E.2 Health and health transition probabilities

To estimate health transition probabilities I first must decide what conditions in the data match with belonging to the various health states in the model. I estimate care need in the data by a multi-step process:

1 - Using the estimation sample, I regress the probability of a parent receiving *any* type of care - be it formal or informal, from any source - on a set of objective health measures in the HRS[54]. Although care receipt is endogenous, the existence of means-tested formal care through Medicaid suggests that care receipt from any source will be a reasonable indicator of actual care need.

[52] 29% of children in the estimation sample with positive care hours reported providing less than 8 hours per month, but these children provided only 1.4% of the aggregate care hours.

[53] To calculate these means, I winsorize care hours at the 10th and 90th percentile to reduce the influence of outliers.

[54] In particular, the explanatory variables are a dummy for 0 ADL difficulties, a dummy for 0 IADL difficulties, quadratics in ADL difficulties and IADL difficulties, all interacted with a dummy for whether the parent suffers from a memory disease and interacted with a quadratic in age.

2 - I use the regression to assign predicted care need to for each parent in each wave, i.e. the probability of receiving any type of care that their objective health measures imply, according to the regression.

3 - I class as having moderate care needs ($h_t^i = 1$) all those between the 90th and 95th percentile of predicted care need and I class as having severe care needs ($h_t^i = 2$) all those who above the 95th percentile

Table 1.15 below presents statistics on objective health measures by assigned h_t^i. As expected, ADL difficulties, IADL difficulties and memory disease are very low for those with $h_t = 0$ and increase with h_t.

Table 1.15: Objective health statistics by model health state

	$h_t = 0$	$h_t = 1$	$h_t = 2$
# ADL difficulties	0.22	2.16	2.98
# IADL difficulties	0.12	2.46	3.81
Memory disease	0.01	0.11	0.39

Notes: h_t=0, 1, 2 refer to no care needs, moderate care needs and severe care needs respectively. Each value in the table is the mean of the left hand variable for people in the relevant health state in the model estimation sample.

Then, for the health transition probabilities themselves I estimate an ordered probit model for period t health state as a function of $t - 1$ health state as well as age, age squared and permanent income, using the model estimation sample.

1.E.3 Work hours and income process

I assume that if a child is working (a state which captures both full- and part-time work) then they work 35 hours per week.

As for the income process, I assume the log of equivalised income is linear in parameters. In particular, $\log(w)$ is a function of age, age squared, gender, couple status, education, work choice in the previous period and work choice in the current period. The reason why income in t might depend on work choice in $t - 1$, even conditional on work choice in t, is that leaving the labour market to provide care will impose a penalty on future wages through, for instance, loss of human capital. This is a mechanism examined by Skira (2015) in assessing the cost of care provision in terms of foregone current and future wages.

The HRS reports child income only in broad brackets so is not very useful for getting at measures of e.g. the gender wage gap. Instead, I use data from the PSID Family File (*PSID* 2024) from 1999 to 2019. I select an estimation sample of individuals with one sibling who are aged between 21 and 60. For this estimation sample, I regress log equivalised income on explanatory variables including demographics and labour market choice. Table 1.16 presents the results of this estimation.

Table 1.16: Estimated income process

	Dependent variable:
	Log of biennial equivalised income
Constant	7.775***
	(0.203)
Age	0.082***
	(0.010)
Age sq.	−0.001***
	(0.0001)
Female	0.107**
	(0.037)
College	0.547***
	(0.029)
Couple	0.372***
	(0.018)
Couple × Female	0.363***
	(0.024)
Worked in t	1.384***
	(0.223)
Age × Worked in t	−0.035***
	(0.010)
Age sq. × Worked in t	0.0003*
	(0.0001)
Worked in $t-1$ × Worked in t	0.518***
	(0.051)
College × Worked in t	−0.049
	(0.032)
Female × Worked in t	−0.234***
	(0.068)
Female × Worked in $t-1$ × Worked in t	−0.182**
	(0.061)
Observations	19,544
Adjusted R^2	0.315

Notes: estimation via OLS. Data from the PSID Family File 1999-2019. Individuals considered to be Working if they work more than 100 hours per year. Household income is equivalised by dividing by the square root of the number of household members, counting children as half a household member. Regression weights are PSID household-level weights. *p<0.05; **p<0.01; ***p<0.001.

1.E.4 Other parameters

To estimate the mean and variance of the wealth shock I consider wealth changes in my estimation sample of HRS data for people who are healthy in both waves, hence who do not face any impact of long-term care costs on their wealth. I winsorize the wealth changes at the 10th and 90th percentiles to reduce the impact of measurement error. The mean and standard deviation of the wealth changes are -\$2.0k and \$115.3k respectively.

A parent's maximum age is assumed to be 100. When parents die their bequest is split equally between their children. As in Ko (2022), the children then consume the bequest over the next T_{beq} periods while working full-time – this provides a terminal payoff for the children to close the model. I assume $T_{beq} = 5$ so children spread their bequest consumption over the following 10 years.

I assume parents in the "older" group (70-85 years old in their first observation in the data) start the model aged 78 and parents in the "younger" group (55-69) start the model aged 63. I assume that the elder child is 24 years younger than the parent, and the younger child is 29 years younger than the parent, matching mean age gaps in the HRS data.

Finally, I assume the discount factor β is equal to 0.93. Note that each period lasts 2 years so this corresponds to an annual discount factor of $\sqrt{0.93} = 0.964$.

1.F Pseudo maximum likelihood estimation

The likelihood function for the observed choices is given by:

$$L^*(\theta) = \prod_{n=1}^{N} \prod_{\tau=1}^{T_n} P^A_{\sigma^*}(d^A_{n\tau}|s_{n\tau}, \theta) P^B_{\sigma^*}(d^B_{n\tau}|s_{n\tau}, \theta) \tag{1.22}$$

where T_n is the total number of periods for which family n with children A and B is observed in the data and $P_{\sigma^*} = \{P^A_{\sigma^*}, P^B_{\sigma^*}\}$ is the set of optimal decision rules for the two children, obtainable through solving the model fully.

To avoid the significant computational cost of solving the model fully, I instead maximise the pseudo likelihood function (Aguirregabiria and Mira 2007). This uses an approximation of P_{σ^*} using the first-stage estimates of the value functions of the two agents. In particular, given a set of first-stage policy function estimates $\hat{\sigma} = \{\hat{\sigma}^A, \hat{\sigma}^B\}$, the pseudo likelihood function will be:

$$L(\theta, \hat{\sigma}) = \prod_{n=1}^{N} \prod_{\tau=1}^{T_n} \Lambda^A_{\sigma^*}(d^A_{n\tau}|s_{n\tau}, \hat{\sigma}, \theta) \Lambda^B_{\sigma^*}(d^B_{n\tau}|s_{n\tau}, \hat{\sigma}, \theta) \tag{1.23}$$

where $\Lambda^i_{\sigma^*}(d^i_{n\tau}|s_{n\tau}, \hat{\sigma}, \theta)$ is the policy iteration operator for child i, given by:

$$\Lambda^i_{\sigma^*}(d^i_{n\tau}|s_{n\tau},\hat{\sigma},\theta) = \frac{\exp(\hat{v}^i(s_{n\tau},d^i_{n\tau},\hat{\sigma},\theta))}{\sum_{d^{i\prime}_{n\tau}\in F^i_{n\tau}}\exp(\hat{v}^i(s_{n\tau},d^{i\prime}_{n\tau},\hat{\sigma},\theta))} \tag{1.24}$$

In other words, the policy iteration operator is an approximation of the true optimal decision rules which updates the first-stage policy function estimates $\hat{\sigma}$ by using these first stage policy function estimates to generate implied choice specific value functions $\hat{v}(\cdot)$ and then using these choice-specific value functions to recover approximate optimal decision rules in each state. The $\hat{v}(\cdot)$ terms are recovered by the simulation procedure described in the main text.

I then maximise the pseudo likelihood to recover the two-step CCP estimator $\hat{\theta}$:

$$\hat{\theta} = \arg\max L(\theta,\hat{\sigma}). \tag{1.25}$$

1.G Evaluating counterfactuals

In this appendix I outline the approach for evaluating counterfactuals in models with multiple equilibria, set out in Aguirregabiria and Ho (2012).

A (Markov Perfect) equilibrium of the game can be written as a fixed point:

$$\mathbf{P} = \Psi(\theta,\mathbf{P}) \tag{1.26}$$

where $\mathbf{P}$ is the vector of choice probabilities for each player in each state, $\Psi(\cdot)$ is a vector-valued best response function for each player and state and θ is the vector of parameters. In other words, in equilibrium, choice probabilities must be best responses to everyone's choice probabilities, given the parameters.

A complication in this model is that there are multiple equilibria: there are multiple solutions to Equation 1.26. The model is thus completed by an equilibrium selection mechanism $\pi(\theta)$, which selects a set of equilibrium choice probabilities from all the possible sets of equilibrium choice probabilities associated with θ.

Let $\mathbf{P_0}$ be the true population choice probabilities and let θ_0 be the true parameter vector governing these choices. It must be the case that $\mathbf{P_0} = \Psi(\theta_0,\mathbf{P_0})$. Even though the exact form of $\pi(\theta)$ is not known, it is known that $\pi(\theta_0) = \mathbf{P_0}$.

Let $\hat{\theta}$ and $\hat{\mathbf{P}}$ be consistent estimates of θ_0 and $\mathbf{P_0}$. Suppose a researcher is interested in what the counterfactual equilibrium is at θ^*, i.e. the researcher wants to evaluate $\pi(\theta^*)$.

Then, assuming that $\pi(\cdot)$ is continuously differentiable around $\hat{\theta}$, the researcher can use the following Taylor approximation to approximate $\pi(\theta^*)$ around $\hat{\theta}$:

$$\pi(\theta^*) \approx \pi(\hat{\theta}) + \frac{d\pi(\hat{\theta})}{d\theta'}(\theta^* - \hat{\theta}) \tag{1.27}$$

Using the fact that $\pi(\hat{\theta})$ is equal to $\hat{\mathbf{P}}$ and to $\Psi(\hat{\theta},\hat{\mathbf{P}})$, one can differentiate $\pi(\hat{\theta})$ with respect to θ and solve for $\frac{d\pi(\hat{\theta})}{d\theta'}$, substituting into Equation 1.27 to arrive at:

$$\pi(\theta^*) \approx \hat{\mathbf{P}} + \left(I - \frac{d\Psi(\hat{\theta}, \hat{\mathbf{P}})}{d\mathbf{P}'}\right)^{-1} \frac{d\Psi(\hat{\theta}, \hat{\mathbf{P}})}{d\theta'}(\theta^* - \hat{\theta}) \tag{1.28}$$

All objects in Equation 1.28 are known to the researcher or in principle calculable from what is known to the researcher. The expression in Equation 1.28 also captures the fact that the counterfactual equilibrium probabilities will depend both on the direct effect of the parameters on the probabilities, captured by $\frac{d\Psi(\hat{\theta}, \hat{\mathbf{P}})}{d\theta'}$, and the indirect strategic effect through the change in other players' equilibrium choices, captured by $\left(I - \frac{d\Psi(\hat{\theta}, \hat{\mathbf{P}})}{d\mathbf{P}'}\right)^{-1}$.

A complication of this approach is that the matrix $\left(I - \frac{d\Psi(\hat{\theta}, \hat{\mathbf{P}})}{d\mathbf{P}'}\right)$ is very large with dimension equal to the number of states multiplied by the number of players multiplied by the number of choices (less 1). In my case calculating and inverting such a matrix would be very costly.

Instead, I calculate a restricted version of the above equation. In a departure from Aguirregabiria and Ho (2012), I impose that the derivative of best responses in state x to the other player's play in state y is non-zero only when $x = y$. This amounts to assuming that in the new equilibrium child i adjusts their strategy taking account of any change in strategy by child j for the current period but is myopic about any changes to j's (or i's) strategy in future periods. I then calculate Equation 1.28 separately for each state. This means that the dimension of the matrix to be inverted in each equation is now only equal to the number of players times the number of choices (less 1), making the problem tractable. From these restricted versions of Equation 1.28 for each state I recover the counterfactual conditional choice probabilities in each state, and thus carry out counterfactual analysis.

1.H Model fit and counterfactuals

1.H.1 Model fit - location

Figure 1.12 shows proportions of children living within 10 miles of their parent, by child gender and parental health status.

The fit is good, though there is not a large amount of variation to match in the real data.

1.H.2 Counterfactual - no gender wage gap

In the main text I consider a counterfactual where I set preference differences between men and women to 0. I argue that the fact that doing so reduces the gender care gap by around four-fifths shows that the gender *wage* gap is not a significant driver of the gender

Figure 1.12: Proportion of children living near parents - data versus model

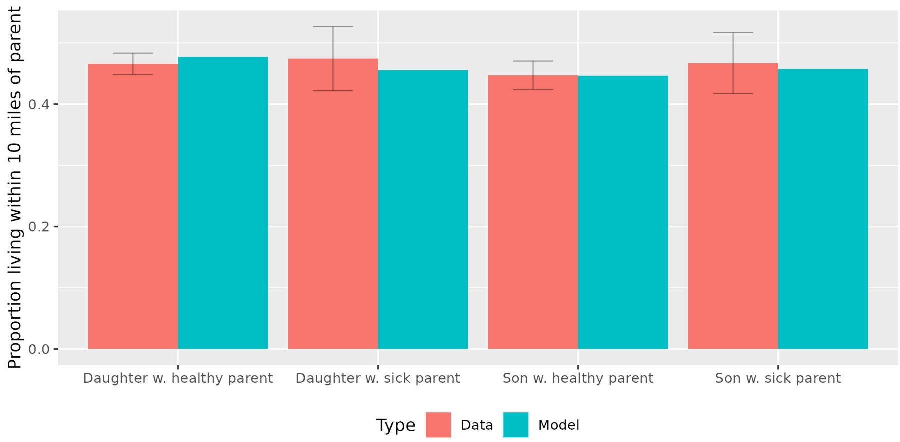

Notes: standard errors clustered at household level. For the four cases, N in the real data is 7832, 761, 7582 and 713 respectively.

care gap. It is not differences in observed opportunity costs between men and women, but rather their differences in preferences[55] that creates the gap in care provision. To illustrate this point further, I conduct an extra counterfactual experiment: I set eliminate the gender wage gap, so that women's equivalised income function is exactly the same as men's in the original model[56].

Figure 1.13 below shows the result of this counterfactual exercise. There is negligble change in the caregiving behaviour of women and men: the gender care gap shrinks by 0.3%. Thus, this backs up the finding in the main text that it is overwhelmingly preference differences which drive the gender care gap.

[55] As discussed above, "preferences" here includes the effects of e.g. social norms placed on women as care providers.

[56] Each child's value function is non-linear in the wage, so it is more difficult to evaluate counterfactuals by taking the derivatives of the probabilities of providing care with respect to the wage parameters and using these as part of the Taylor approximation approach set out in Appendix 1.G. For this reason I do not use the full method as set out in Appendix 1.G for this counterfactual. Instead, I simply alter the wage gap parameters, holding all other parameters constant, and recalculate agents' value functions and resulting behaviour, ignoring any complications around multiple equilibria. For this reason, the counterfactual results presented here are less reliable than those in the main text (though they are broadly consistent with those in the main text).

Figure 1.13: Counterfactual - no gender wage gap

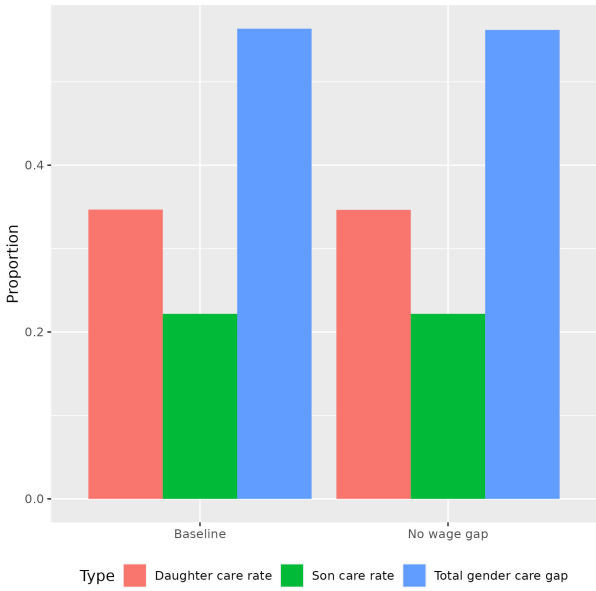

Notes: “care rate” is the probability of providing care conditional on having a parent with care needs. “Total gender care gap” how much larger, in percentage terms, daughters’ care rate is relative to sons’ care rate.

2 | Where's my pension gone? Labour supply effects of mistakes in retirement planning

Where's my pension gone? Labour supply effects of mistakes in retirement planning

Abstract

I analyse how agents respond to realising mistakes in their retirement planning using a reform to the state pension age for women in the UK. This reform raised the age at which women could start receiving their public pension by several years. However, many women were not aware that this reform had been put in place, leading them to have mistaken beliefs about when they would be able to start receiving their pension. I first document the correlates of mistaken beliefs using data from the English Longitudinal Study of Ageing (ELSA) and present descriptive evidence suggesting an increase in labour supply after people realise they were mistaken about their pensions. I then build a life-cycle model of consumption and labour supply around retirement, where agents face a cost of cognitive effort of paying attention to reforms to the pension system, and I estimate the model by matching data moments from ELSA for the cohorts of affected women. I find that mistaken beliefs about the reform cost inattentive households £457 on average, though with notable heterogeneity by demographics, and that agents in the model significantly increase labour supply once they realise their mistakes.

2.1 Introduction

A longstanding concern for policymakers and researchers is whether people are saving adequately for retirement (Banks, Emmerson, et al. 2005; Crawford and O'Dea 2020; Scholz, Seshadri, and Khitatrakun 2006). In particular, over-optimism about medical costs (Department of Health 2011), income in retirement (Hentall-MacCuish 2025) or underestimating longevity (O'Dea and Sturrock 2018) could lead to people having inadequate resources and hence a lower quality of life in retirement. Governments around the world have taken action to overcome these behavioural biases such as auto-enrollment in private pension schemes (OECD 2019).

The existing literature has taken several approaches to assessing whether people are saving adequately for retirement. Some authors have tried to assess whether consumption drops around retirement are evidence of drops in welfare from bad planning (Battistin et al. 2009; Olafsson and Pagel 2018). Others have taken a more structural approach, building models of agents' decision-making around retirement and assessing the extent to which heterogeneity in preferences (Crawford and O'Dea 2020) or beliefs (O'Dea and Sturrock 2020) can explain behaviour which appears to be evidence of bad planning.

A key difficulty with analysing the issue of optimal saving for retirement is that it is rare that we are able to identify who is planning badly for retirement. In principle, someone might have low savings in their late working career because of previous negative shocks they have suffered, or because they discount the future very strongly, or because they have resources which are not observed by the econometrician. In all these cases it is difficult to argue that these people are making mistakes, or are planning badly. To understand whether people are making mistakes, we need to have some information on the beliefs they have about the future: if they have drastically incorrect beliefs about how much income they will have in future, or how likely they are to face crippling medical costs, then we can be much more confident in saying that they are making mistakes. Moreover, to understand the extent to which these planning mistakes are detrimental to welfare, we must be able to observe how agents change their behaviour - and the costs they incur from doing so - when they realise their mistakes. For instance, if people had highly inaccurate beliefs about their future pension income, but their behaviour would be similar whether or not their beliefs were accurate, then their planning mistakes may not be very costly.

This paper tackles this difficulty by using household survey data on subjective beliefs about future pension income eligibility, where these beliefs can be easily verified as being correct or not. This allows me to identify i) who is making mistakes, ii) exactly how the agents are misestimating their income in retirement and iii) how agents respond after they realise their mistakes. As such, the contribution of this paper is to use a new source of data to gain a cleaner and more direct insight into the prevalence, drivers and consequences of inadequate planning for retirement.

More specifically, this paper examines these issues by exploiting a reform to the state pension age (SPA) - the earliest age at which people can claim a public pension - for women in the UK. This reform was originally announced in the 1995 Pensions Act, 15 years before it came into force. For certain cohorts of women, their SPA increased by several years, from 60 to (for the most affected women) 65. Given that in my data state pension income makes up around 64% of the median single retiree household's income and around 47% of the median couple retiree household's income, this amounted to a very substantial change in income for women looking ahead to their early 60s.

The key feature of the reform, for the purposes of this paper, was that many women were unaware of how the reform affected them, even when close to retirement (Holman, Foster, and Hess 2020): only 43% of affected women knew their SPA had increased in 2004, 9 years after the passage of the 1995 Pensions Act and 6 years before its provisions would come into force (PHSO 2024). As such, many will have had a "moment of realisation" shortly before their original planned retirement date when they realise that a future source of income they were relying on will instead be delayed by several years. Using high-quality panel data from the English Longitudinal Study of Ageing I am able to track the scale and nature of these mistakes, and how agents responded once they realised their mistakes.

An important advantage of this setting for analysing the (in)adequacy of agents' planning for retirement is the simplicity of the object - the SPA - that agents have beliefs over. For instance, if we observe an agent with incorrect beliefs about their SPA, we can assess by how much they are misestimating their income in retirement, and once they realise their mistake, this is plausibly a relatively clean shock to expected future income. In contrast, if someone suffers a health shock, then this reduces their expected future income, but the amount of the reduction is unclear and any behavioural responses to the change in their future income will be conditioned by their new health problem, rather than being just a response to changes in expected future income.

However, a complication is that the correctness of agents' beliefs about the SPA is endogenous. Agents for whom the "stakes" (Brown and Jeon 2024) of the issue of having correct pension beliefs are high - e.g. those whose optimal behaviour is highly sensitive to what their expected future income is - will plausibly be more likely to do their research to ensure that they know the truth of their SPA. Moreover, steps people might take to protect themselves against drops in future income, such as re-entering the labour market, are also plausibly steps that make it more likely for people to acquire information about the pension system. Thus, understanding how people respond to acquiring better information about the adequacy of their retirement plans - and thus understanding the extent to which people plan adequately for retirement - will require a model of endogenous information acquisition. To this end, I develop and estimate a rich quantitative life-cycle model of labour supply and consumption around retirement with rational inattention about the SPA and use this to assess agents' responses to learning about the adequacy of their retirement plans.

The paper proceeds as follows. First, I use ELSA data on women's beliefs about their SPA to identify mistaken beliefs as those where what the respondent believes to be their SPA is sufficiently far from the true SPA and I examine the correlates of mistaken beliefs in the data. I find that homeowners, those in the labour force, those with private pensions and those who have a long financial planning horizon are more likely to have correct beliefs, though strikingly women in couples and with a college degree are actually no more likely to have correct beliefs. This is consistent with the agents with higher education or in couples facing both lower costs of acquiring information - due to plausibly higher financial literacy and having an extra source of information in the household, respectively - and facing lower incentives to acquire information, because the state pension is a smaller fraction of their expected income in retirement. Moreover, I find some evidence that when agents transition from having incorrect beliefs to having correct beliefs in the data both their and their spouses' probability of working increases, as does their annual earnings, though there is no obvious corresponding change in consumption, asset holdings or measures of financial well-being. While the endogenous nature of information acquisition prevents a simple causal interpretation, these findings are consistent with agents cushioning the blow of lower expected future income by working more rather than reducing their consumption or drawing down their assets.

Then, I develop the quantitative model which is the focus of the paper. Agents start the model at age 50. They are either attentive of their true SPA, or inattentive. For attentive agents, the model consists of a labour and consumption choice every period, where agents face different wages and costs of working according to their demographics and state variables. For inattentive agents, there is an additional choice every period of whether to "do their research" about the true SPA and become attentive, where the costs of becoming attentive also vary by demographics and state variables.

I estimate a large number of parameters within the model, allowing for significant heterogeneity in wages, costs of working and costs of becoming attentive. The model is estimated by the Method of Simulated Moments, matching data moments on women and men's labour supply, labour earnings, correctness of beliefs and responses to realising mistakes. Notably, I estimate that agents face significant costs of re-entering the labour force in terms of wages and disutility of working, exacerbating the impact of inadequate planning for retirement, and that the costs of becoming attentive vary importantly in the population, with agents in couples, with college educations and attached to the labour force facing lower costs.

Using the estimated model I find that agents indeed increase their labour supply and work longer into their 60s once they realise that their SPA beliefs were in error. Moreover, I quantify the extent of losses because of inattentiveness to the reform: on average ex post, households who start the model with incorrect beliefs would be indifferent between a £457 increase in wealth at the start of the model and living in a world where they had always had correct beliefs. This cost is modest relative to a median household consumption at age 60 in the model of £20.3k per two-year period. By decomposing the costs of the imperfect communication of the reform I find that cohorts who faced smaller changes in their SPA, or who were further from retirement at the start of the model, faced less cost, as did those with more education who were protected against the consequences of mistakes.

This paper contributes to two different literatures. Most narrowly, it adds to work on the effects of changing statutory retirement or pension eligibility ages across several countries. This type of policy reform has been studied in the context of the UK (Cribb, Emmerson, and Tetlow 2016), US (Deshpande, Fadlon, and Gray 2024), Austria (Staubli and Zweimüller 2013), Denmark (García-Miralles and Leganza 2024a) and Japan (Nakazawa 2022), amongst others. The literature has generally found that people adjust their labour supply to work longer as the SPA moves back (García-Miralles and Leganza 2024b). The contribution of this paper relative to other work in this area is the exploiting of agents' beliefs and mistakes about their SPA, rather than just using the change in the statutory early retirement age. By doing so, I can use more variation in the timing and size of expected income shocks to assess the effects on labour supply, and I can capture the fact that ignorance of the nature of reforms can modify their welfare impact.

The paper is closest in topic and approach to contemporaneous work by Hentall-MacCuish (2025), who similarly studies UK women's SPA mistakes using a dynamic model of rational inattention. He finds that costly attention and ignorance of the true SPA can explain an empirical puzzle where people are disproportionately more likely to retire at exactly their SPA even though the financial incentives to do so are apparently limited, building on previous work in this area by Cribb, Emmerson, and Tetlow (2016). While the current paper and Hentall-MacCuish (2025) share a policy context and broad empirical approach, there are important differences in model construction. Hentall-MacCuish (2025) develops a sophisticated model of dynamic rational inattention whereby agents continually update their beliefs about the SPA every period depending on previous information; this allows detailed consideration of the effects of rational inattention on labour supply but requires simplifications elsewhere in the model, such as imposing inelastic labour supply for spouses, estimating wage processes outside of the model and cutting back on state variables that could determine incentives to work, such as previous labour supply. In contrast, the current paper offers a much simpler model of rational inattention, where agents are either attentive or inattentive. This implies a less detailed treatment of the drivers of incorrect beliefs but allows more complexity in parts of the model that are relevant to assessing the welfare effects of mistakes in SPA beliefs. For instance, in the current paper, both women and men in the household make labour supply decisions and there is much more heterogeneity in wages, disutility of working and costs of paying attention by a broad set of exogenous and endogenous state variables. As such, this paper is able to contribute a detailed assessment of the welfare effects of mistakes in retirement planning.

A second, broader, literature to which this paper contributes is the large literature featuring life-cycle models of labour supply and consumption around retirement, particularly those concerning the adequacy of preparations for retirement. The literature on labour supply around retirement is surveyed in Blundell, French, and Tetlow (2016). More general life-cycle papers on behaviour around and into retirement focussed on precautionary saving (De Nardi, French, and John Bailey Jones 2016), medical costs (De Nardi, French, and John B. Jones 2010) and bequest motives (Lockwood 2018), while other authors have assessed the importance of heterogeneity in preferences (Crawford and O'Dea 2020) or beliefs (O'Dea and Sturrock 2020). The current paper's focus on quantifiable and observable mistakes in beliefs about future income allows me to contribute an analysis of how people tend to respond to realising their mistakes in retirement planning and what determines the welfare costs they suffer when doing so.

The rest of the paper proceeds as follows: in Section 2.2, I establish some basic empirical facts about the pension reform and mistaken beliefs. In Section 2.3, I present the quantitative model. Section 2.4 discusses estimation, identification and model fit, while Section 2.5 contains welfare analysis and counterfactuals. Section 2.6 concludes.

2.2 Empirical facts

2.2.1 Policy context

The UK state pension is a regular payment from the government that can be claimed from an individual's state pension age (SPA) onwards. The amount received depends on the years of National Insurance (NI) contributions that an individual has made, with those making the maximum amount of NI contributions receiving a "full" state pension. There is also a small earnings-based component which applied differently for different cohorts though as Figure 2.1 below suggests in practical terms the absolute amount received did not vary much along the income distribution. In 2025-26, the "full" state pension is £230.35 per week, or approximately £12k per year (Department for Work and Pensions 2025a). In Appendix 2.A.1 I discuss in more detail the exact structure of the state pension.

The state pension is an important source of income for retired households. Figure 2.1 below shows a breakdown of income by source across the income distribution for retiree households in 2018, the most recent wave of ELSA data that I use. For the median single household, 64% of their income comes from public pensions[1], whereas for couple households, this figure is 47%. There is little variation in the absolute amount of the state pension across the income distribution, with instead increases in private pension being notable for richer households relative to poorer households.

Figure 2.1: Income distribution by quintile - retirees

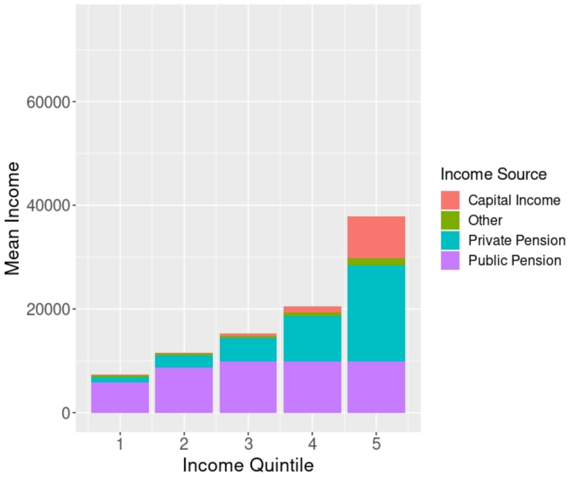

(a) Singles

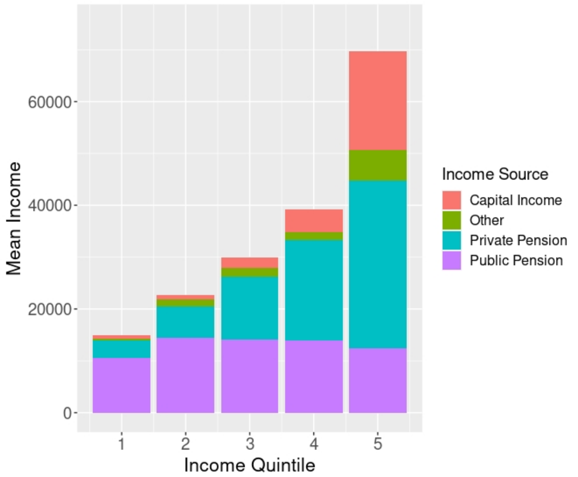

(b) Couples

Notes: single households where the household member self-describes as retired, and couple households where both household members self-describe as retired. Quintiles weighted by household-level weights (means of ELSA person-level weights within household).

2.2.1.1 Reform to women's SPA

From the 1940s to 1995, the SPA had been 65 for men and 60 for women (Mackley, Thurley, and McInnes 2021). In 1995, the Pensions Act was passed, decreeing that the women's SPA would rise from 60 to 65 gradually between 2010 and 2020. Thus, women born after (April) 1950 - who were 45 or younger at the time of the reform - would receive their state pensions later than the age of 60. In 2011, a new Pensions Act was passed, accelerating the rise in the SPA. Figure 2.2 below shows SPA for women, before the 1995 Pensions Act, under the 1995 Pensions Act and under the 2011 Pensions Act.

The left-hand graph shows the prevailing SPA for women in a particular calendar year, under the pre-1995 regime, the 1995 Pensions Act and the 2011 Pensions Act. The 1995 Pensions Act only affected SPAs from 2010 onwards, i.e. 15 years after the passage of the Act, and set out a plan for the SPA for women to rise to 65, i.e. equality with men, by 2020. The 2011 Pensions Act accelerated the increase to 65 between April 2016

[1] In the figure, "public pensions" groups together the state pension and other associated public pensions such as widow's pension and disability pensions. On average, the state pension's share of total public pension income in the data is 95%.

Figure 2.2: SPA for women by year of receipt and by year of birth

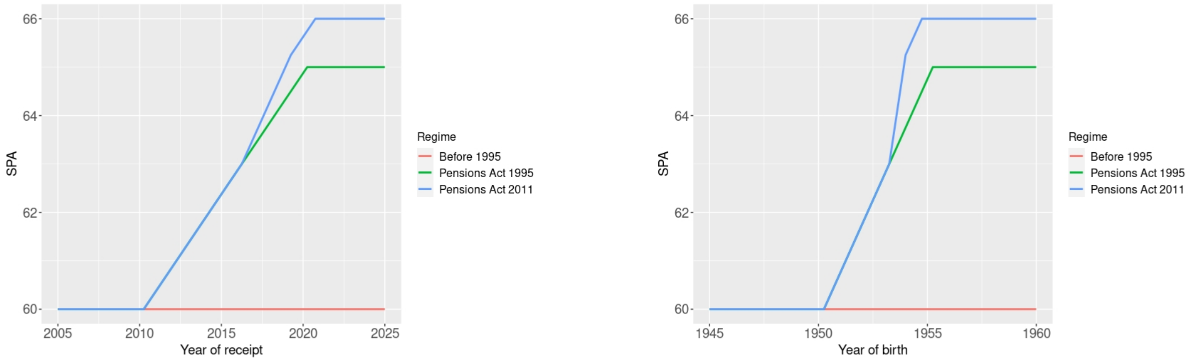

(a) By year of first receipt (b) By year of birth

Notes: Data from Department for Work and Pensions (2014).

and November 2018, and then added an extra increase up to 66 (for men and women) between November 2018 and October 2020.

The right-hand graph shows the SPA for different women by year of birth instead. Only those born after 1950 were affected, with more significant changes for younger cohorts. The 2011 Pensions Act brought particularly large changes for those born in 1953 or later, who would have been 58 or younger at the time of the reform.

As Mackley, Thurley, and McInnes (2021) describe, a key aspect of the reform was that there was widespread ignorance about the consequences among affected groups. The Parliamentary and Health Service Ombudsman (PHSO) found in 2020 that maladministration had taken place in the communication of the reform to women whose SPA was increasing. Department of Work and Pensions research indicated as early as 2004 that there was widespread misunderstanding of the reform, and the Department proposed in November 2006 that it write directly to affected women to inform them of the changes, but the PHSO found that the letters were not sent until March 2009 (PHSO 2021). The PHSO then subsequently found in 2024 that this maladministration had led to injustice and proposed that the government pay out in compensation to affected women (PHSO 2024), a proposal that has been rejected at the time of writing (Department for Work and Pensions 2024). It is this miscommunication and misunderstanding of the reform that allow me to assess the labour supply and welfare consequences of mistakes in retirement planning.

2.2.2 Descriptive evidence

2.2.2.1 Labour supply by cohort

To establish some descriptive facts, I use data from the English Longitudinal Study of Ageing (ELSA), a representative panel survey of the non-institutionalised over-50 population of England. Panel respondents are asked a battery of questions about their demographics, wealth, income, activities, spending and pensions. I use ELSA data from 2004 to 2018 comprising 8 biennial waves (Waves 2 to 9). Appendix 2.B.1 presents some descriptive statistics on the estimation sample.

To understand basic trends in labour supply I first consider labour force participation by age across cohorts with different SPAs. Figure 2.3 below plots the proportion of individuals in the labour force (i.e. working full-time, working part-time, or unemployed) by age across different cohorts, separately for men and women.

Figure 2.3: Proportion in labour force by birth cohort

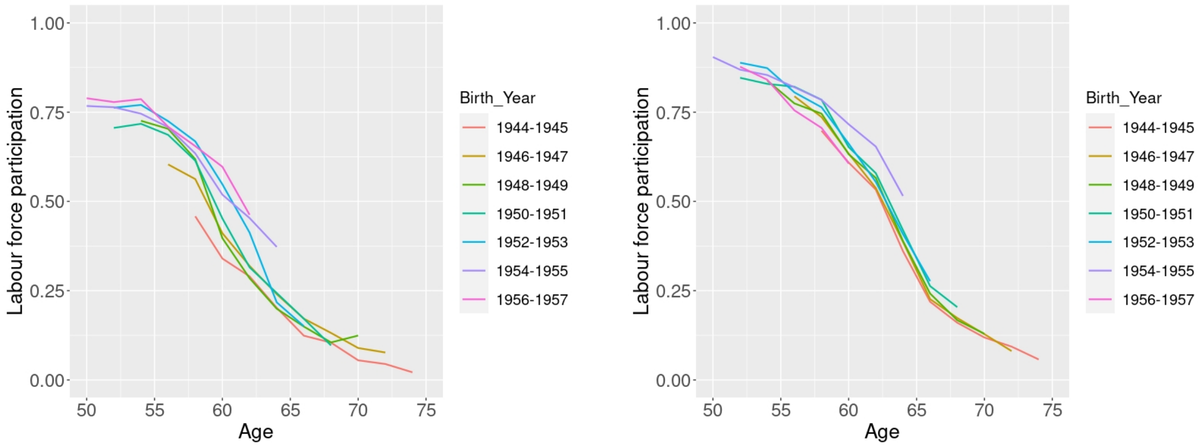

(a) Women (b) Men

Notes: individuals counted as being in the labour force if self-describe as working full-time, working part-time, or unemployed. Means weighted by ELSA person-level weights.

Considering first the left-hand graph for women's labour supply, for pre-reform cohorts - i.e. women born before 1950 - the employment profile by age looks similar, with a particularly sharp drop off at 60, the SPA for these cohorts[2]. However, for later cohorts, women seem to stay in the labour force longer. This suggests that the first-order effects of the reform are to delay women's retirement, in line with the evidence in Cribb, Emmerson, and Tetlow (2016) and Hentall-MacCuish (2025). For men in the right-hand graph, there are no clear cohort effects and the profile of labour force participation continues similarly across cohorts, as is consistent with there being no notable changes to men's SPA[3].

[2]The question of why there is so much bunching in retirement around the SPA is interesting and important. Cribb, Emmerson, and Tetlow (2016) and Hentall-MacCuish (2025) consider this question in detail in the UK context, with the former emphasising the importance of SPA receipt as a signal of when it is appropriate to retire, and the latter emphasising the importance of learning about the SPA.

[3]It is possible that part of the change in women's labour force participation over their 50s and 60s following a change in the SPA is mechanical, e.g. if employment contracts and normal retirement ages are directly linked to the SPA. I consider this possibility in Appendix 2.B.2 by analysing ELSA data on the normal retirement ages associated with respondents' pensions. I show that although the average normal retirement age is increasing over time there is little evidence of widespread mechanical links with the SPA.

2.2.2.2 Prevalence and nature of mistakes

From 2006 (Wave 3), ELSA asked women respondents below the SPA two questions about their perception of the SPA. Respondents were asked "Do you know what age in years and months you will reach the SPA?" (and were then prompted to supply the details). Respondents were also asked "Do you know that the SPA for women is changing?".

Figure 2.4 plots people's responses as to their SPA, separately for different cohorts. The data are taken from Wave 3 of ELSA. In each case, the dotted red line shows the median true SPA for this cohort[4].

The 1949 cohort - the final pre-reform cohort - has SPA beliefs that are tightly bunched on the true SPA of 60. In subsequent cohorts, as the SPA increases (the red line shifts rightwards), there are three main types of respondents. In order of size, these are those who shift in line with the truth, those who still believe that the true SPA is 60, and those who overcorrect to believing that their SPA is higher than the true value, typically believing they have a SPA of 65.

As for how the correctness of beliefs changed over time, Figure 2.5 plots the proportion of people with correct beliefs[5] over time. In particular, I plot 3 lines: the red line is the proportion of women who were aware of the reform to the SPA (in the sense that they answered in the affirmative to the question "Do you know that the SPA for women is changing?"). The green line is the proportion of people with correct beliefs conditional on being aware of the reform, and the blue line is the proportion of people with correct beliefs conditional on not being aware of the reform. The sample from which the means are calculated is women who are below their SPA and whose true SPA is > 62. This second proviso is to ensure that women who continue to answer that they believe their SPA to be 60 even if it is actually higher will be deemed to have incorrect beliefs, given that I am defining "correct" beliefs to be those within two years of the truth.

One clear pattern that emerges is that awareness of the reform in general terms was always very high, even back in 2006. There seems to have been a spike in awareness around 2010, plausibly because this was the first year that people actually experienced not being able to claim their state pension at 60. Correctness of beliefs (whether conditional on being aware or unaware of the reform) increases gradually over time.

However, it is striking that even among those who claim to be aware of the reform

[4] Appendix 2.A.2 discusses how the true SPA measure for each individual is constructed.

[5] In this case, I count a person's beliefs as being correct if they are less than 2 years out from the truth. So, if a person's true SPA is 62 and 3 months, then any answer from 60 and 3 months to 64 and 3 months is deemed as correct. This is obviously a generous definition of what it is to be correct. There are two reasons for this: firstly, in the public ELSA data only the respondent's year of birth, not date of birth, is provided. I impute respondents' month of birth, hence SPA, using a procedure set out in Appendix 2.A.2, but to limit the importance of this imputation procedure to determining the results I focus on "big" mistakes which are clearly the result of the agent's SPA beliefs being incorrect rather than due to their imputed SPA being far from their true SPA. Similarly, by focussing on "big" mistakes, I set aside differences between expected SPA and true SPA which are plausibly due to things like rounding one's SPA to the nearest year or 6 months.

Figure 2.4: SPA beliefs by cohort

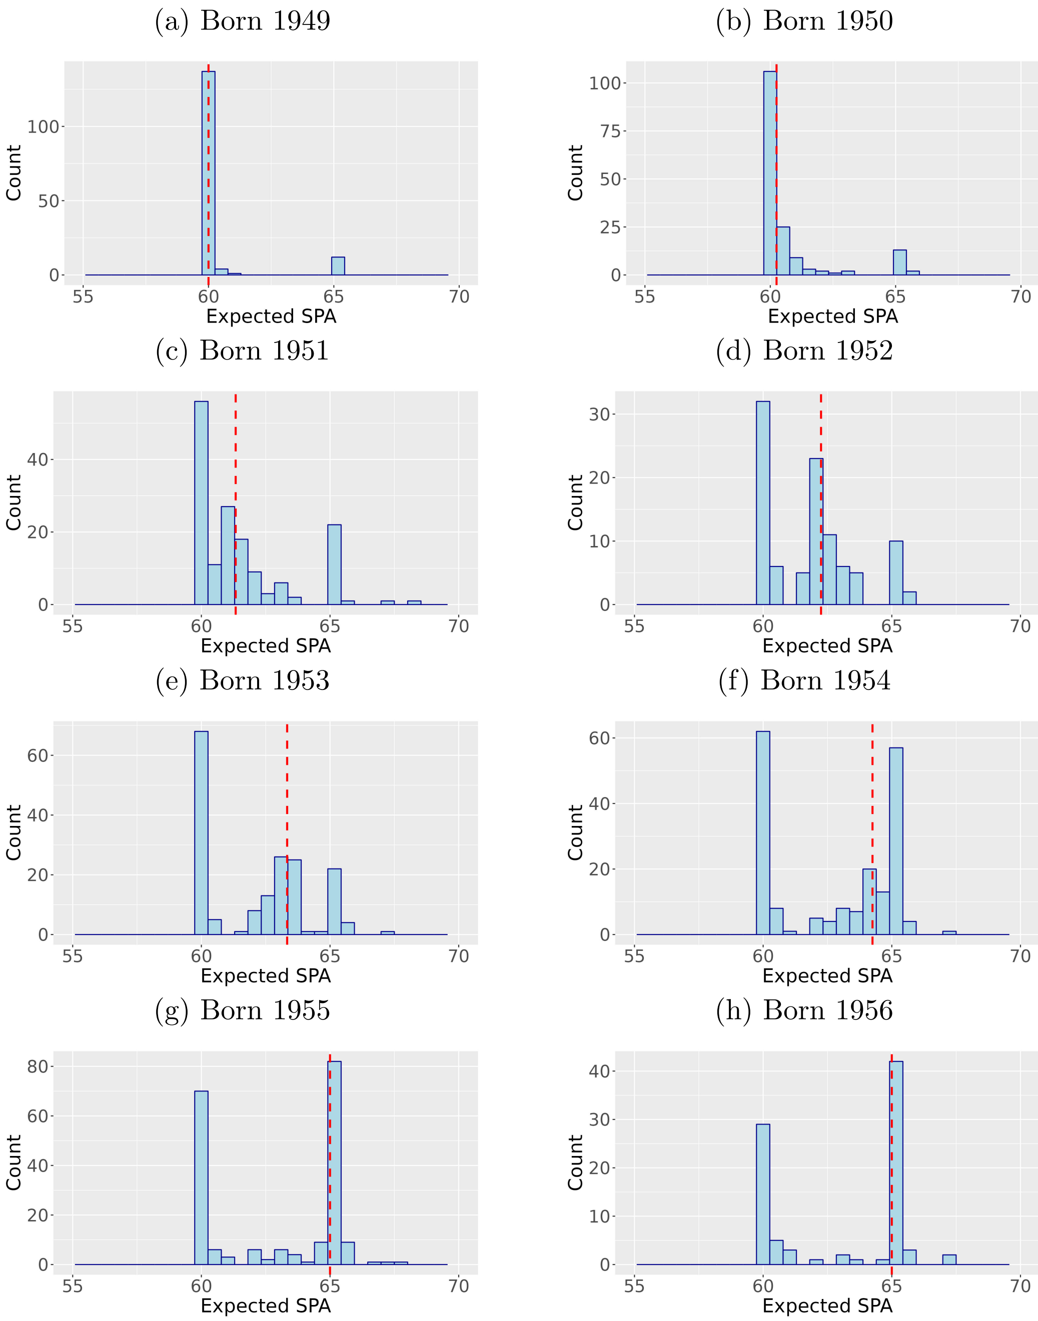

Notes: data taken from ELSA Waves 3 (2006-2007). In each case, dotted red line represents median true SPA for this cohort under the regime in place at the time. Answers of a SPA > 70 or < 55 are omitted - these make up only 1% of answers.

Figure 2.5: Awareness of reform over time

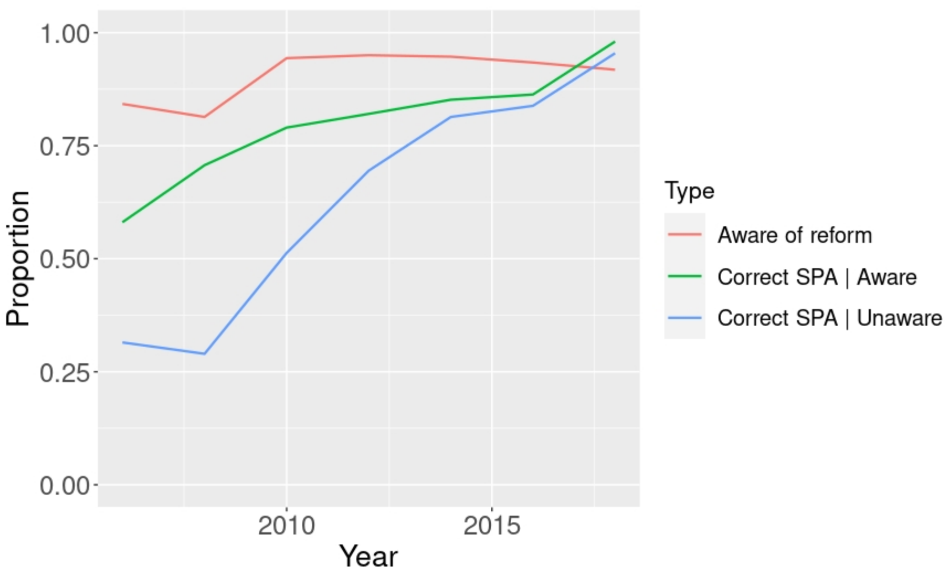

Notes: means calculated from those respondents who are below the SPA and whose SPA > 62. Means weighted by ELSA person-level weights.

in general terms, a significant minority still hold incorrect beliefs about their SPA. For instance, among those who in 2006 claim to know that the SPA for women is increasing, still 42% have incorrect beliefs about their SPA. This suggests that it is not plausible to treat one's realisation about one's true SPA as an exogenous shock. Instead, most people were aware in general that the reform was taking place, and it is those who chose to do further research who were able to find out their true SPA.

Finally, to shed some light on the determinants of having correct beliefs, I regress a dummy for having correct beliefs on a set of RHS variables for Wave 3 (2006) for women below the SPA in that wave and who had data for all the regressors. The coefficients from the regression are displayed in Figure 2.6, with the full regression table in Appendix 2.B.3.

Homeowners, those with private pensions and those who described themselves as being long-term planners[6] were more likely to have correct beliefs, as were those in the labour force, though in this latter case the coefficient is marginally statistically insignificant at the 5% level. There is no strong association between being the ELSA survey financial respondent, or claiming to be socially isolated[7], and having correct beliefs. Finally, those in couples or with a college degree are less likely to have correct beliefs, though again this is not statistically significant at the 5% level.

[6]In Wave 2, ELSA respondents were asked "In deciding how much of your income to spend or save, people are likely to think about different financial planning periods. In planning your saving and spending, which of the following time periods is more important to you?". I class as long-term planners all those who answered that their financial planning period was 5 years or above.

[7]ELSA respondents are asked "How often do you feel isolated from others?". I classify as socially isolated all those who answer "Some of the time" or "Often".

Figure 2.6: Coefficients from correct beliefs regression

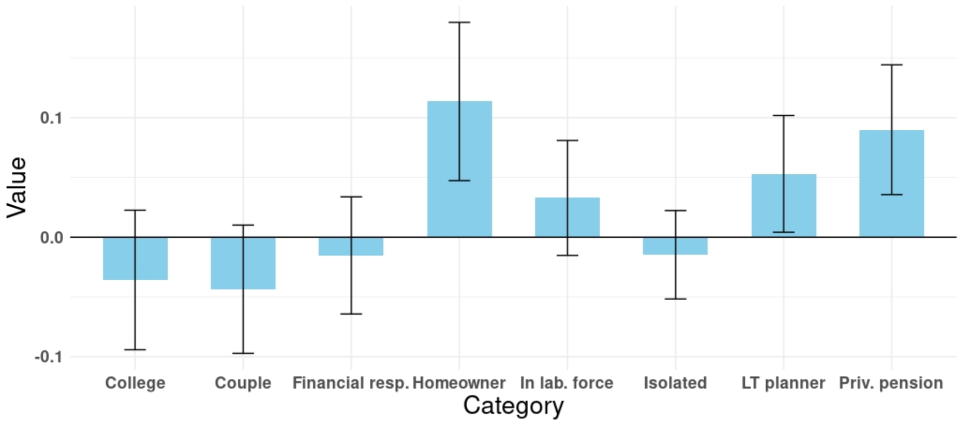

Notes: OLS coefficients from a regression of a dummy for having correct beliefs on the set of RHS variables listed as well as age and YOB fixed effects. Data from ELSA Wave 3 (2006). Confidence bands show 95% confidence intervals. The full regression table is presented in Appendix 2.B.3.

It is somewhat surprising that more educated people, and those in couples, are no more likely, and plausibly are less likely, to have correct beliefs. One way of understanding these results is that while higher-educated people are plausibly more likely to be financially literate, they are also likely to be less reliant on the state pension to support them in old age, so there is less of an incentive to keep track of when exactly their state pension will arrive. Similarly, while couples might be expected to receive more information because of having an extra pair of ears and eyes in the household, the woman's state pension makes up a smaller proportion of total household income in retirement compared to single households.

2.2.3 Responses to realising mistakes

As the ELSA dataset follows individuals over time, it is possible to track how individuals' behaviour changes after they switch from having incorrect beliefs to having correct beliefs. I exploit this information to estimate a set of event studies on choice variables of interest.

Of course, it is difficult to interpret these switches as being exogenous information shocks: people endogenously decide whether to investigate their pension status. Indeed, the evidence from Figure 2.5 suggests that many more people knew that a reform in general was taking place than knew exactly how the reform would apply to them which could be interpreted as many people deciding not to pursue more information on how the reform applied to them. This suggests that there would be selection into those who did the research and therefore received the information. Also, certain choices agents make like re-entering the labour market might expose them to more information, which would be difficult to distinguish from people acquiring information and then entering the labour market as they change their plans. As such, these event studies should be interpreted as giving a description of how choice variables changed around the time of someone switching from having incorrect to correct beliefs, rather than establishing a causal relationship.

For each choice variable of interest y_{it} of individual i at time t, I estimate the following regression:

$$y_{it} = \sum_{j=\{-3,..,3\}} \beta_j D_{t-j} + \alpha_i + \gamma_{at} + \delta_a \times I(\text{In LF in '06})_i + \lambda_a \times I(\text{Spouse in LF in '06})_i \quad (2.1)$$

The event is realising one's mistake, i.e. the first time i is observed with incorrect beliefs in $t-1$ and correct beliefs in t, then i is considered to have realised their mistake at t[8].

Also included as regressors are individual fixed effects α_i, age by time fixed effects γ_{at}, and interactions between age dummies and dummies for whether the individual/their spouse (if any) was in the labour force in 2006, the first observation in the data. Note that the age by time fixed effects will control for people being above or below either the old or new SPA for their particular cohort. The reason for including differential age trends by whether an agent or their spouse was originally in the labour force in 2006 is that we might expect different slopes in the relevant outcome variables by age depending on degree of attachment to the labour force[9]. In Appendix 2.B.4.7 I present the corresponding figures for this section but without these trends by original labour force attachment; the "jump" around the point of realisation is similar but there are pre- and post-trends.

I estimate the model of Equation 2.1 for a range of outcome variables. In each case, the subsample I use is data from Wave 3 (2006) onwards for women born between 1945 and 1956 and who appear in Wave 3, meaning I capture 5 one-year pre-reform cohorts as well as all the cohorts whose SPA was on the transition path between 60 and 65.

Figure 2.7 plots the estimated coefficients from the regression in Equation 2.1 where

[8] There are some instances of agents switching from having correct beliefs to having incorrect beliefs. In my estimation sample, for those below the SPA at time t, the probability of having correct beliefs at t is 0.90 conditional on having correct beliefs at $t-1$, so in around 10% of cases there is a switch from having correct to incorrect beliefs. The probability of switching from incorrect to correct beliefs is 0.57 for those below the SPA. Given these facts, the parameter estimates here should be interpreted as capturing the association between realising one's belief and the outcome variable, including any subsequent changes in one's belief. In Appendix 2.B.4.5, I prevent some supplementary event studies where an agent is counted as realising their mistake in t if they have incorrect beliefs in $t-1$ (and all subsequent periods), correct beliefs at t and also correct beliefs at $t+1$, i.e. counting only those whose switches to having correct beliefs is in some sense stable. Qualitatively, the results are similar.

[9] For instance, suppose there are two groups of people: those more attached to the labour force and those less attached. Suppose that as they get older the probability of both groups working decreases by the same percentage amount every year. This means that the absolute difference in working probability between the two groups would mechanically shrink, creating a trend in the absolute difference in working probability between the two groups that imposing a common age trend would not eliminate. Instead, I impose separate age trends by original labour force status as a proxy for attachment to the labour force.

the outcome variable is a dummy for labour force participation on the LHS. The results suggest that there is an uptick in labour force participation for women after they realise the truth of their SPA, though the standard errors are large.

Figure 2.7: Event study: labour force participation

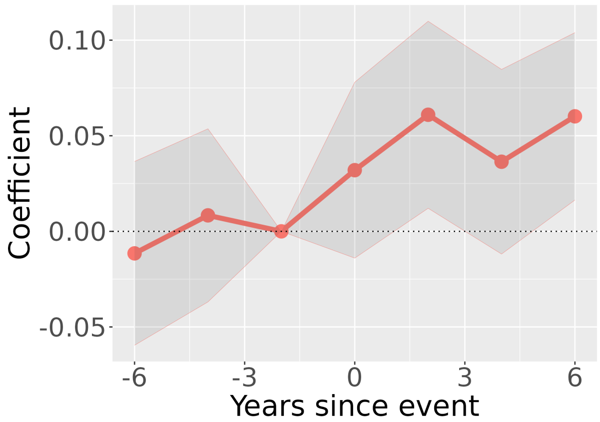

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

In order to understand the mechanism behind this, I use as an outcome variable a dummy for whether an individual switches from being in the labour force at $t-1$ to being out of the labour force at t, i.e. a dummy for leaving the labour force. The results are presented in Figure 2.8.

Figure 2.8: Event study: proportion leaving labour force

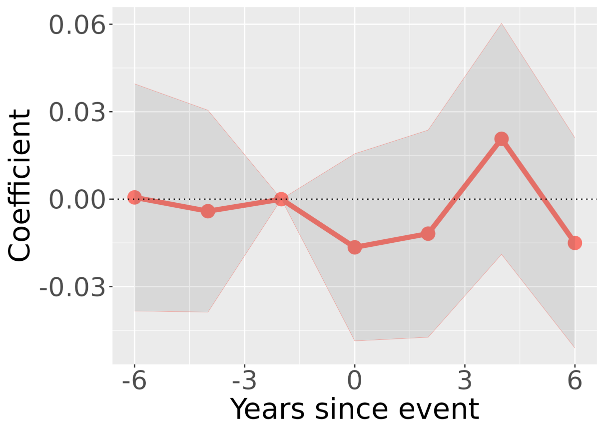

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Again, standard errors are large here so results are not statistically significant, but the figure suggests that one driver of the changes in labour force participation may be that agents who realise their mistakes tend to delay their retirement relative to others: the difference in the rate of leaving the labour force is initially negative before rising to be positive four years after the mistake was realised[10].

[10] In Appendix 2.B.4.4 I show the corresponding graph where the outcome variable is a dummy for

To illustrate the underlying movements in the raw means, Figure 2.9 below plots mean women's labour force participation by age separately for those who had correct pension beliefs at a particular age versus those who did not have correct beliefs at that age. The left-hand panel presents the case where the relevant age is 54-55 and the right-hand panel presents the case where the relevant age is 56-57, with two-year age bins being used because of the biennial nature of ELSA. These ages are chosen because 56-57 is the median two-year age bin at which agents correct their beliefs below the SPA in the sample. In each case, the graph plots mean labour force participation for women, with the only controls being year-of-birth dummies to prevent the graphs simply reflecting compositional changes in the sample across waves[11].

Figure 2.9: Labour force participation by awareness at different ages

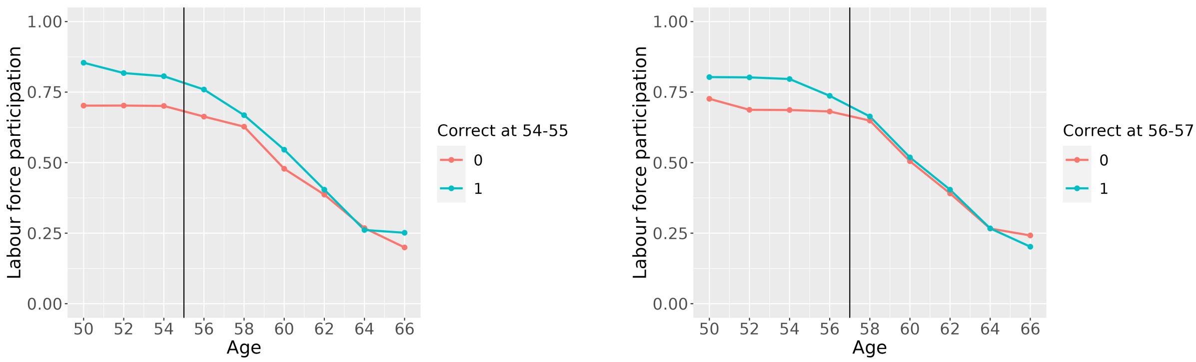

(a) Correct at age 54-55 (b) Correct at age 56-57

Notes: individuals counted as being in the labour force if self-describe as working full-time, working part-time, or unemployed. Vertical black lines indicate the reference age. Only individuals in the sample at the relevant reference age are included. Coefficients plotted are net of year-of-birth controls, using the intercept for the 1952 birth cohort. Means weighted by ELSA person-level weights.

In both panels, the gap in labour force participation is initially substantial but subsequently narrows, indicating that descriptively those who were unaware in their mid-to-late 50s decreased their labour force participation more slowly than those who were aware. In Appendix 2.B.4, I show corresponding figures in terms of labour force participation over calendar time rather than age, and I show in a set of regressions that even controlling for observables having correct beliefs in 2006 is associated with retiring earlier.

As for other outcomes of interest, Figure 2.10 plots the estimates where the LHS variable is the individual's annual earnings. There seems to be an uptick in earnings from labour around when agents realise the truth about their SPA.

rejoining the labour force. In this case, there is no change in pretrends around the point of realising mistakes, suggesting that the main results are more likely to be driven by people delaying leaving the labour force than rejoining the labour force when they realise their mistake.

[11] Specifically, I regress a dummy for labour force participation on age dummies and wave dummies for those with correct beliefs and those with incorrect beliefs at the relevant ages, and plot the coefficients on the age dummies added to the intercept for the 1952 cohort, the median cohort in the sample.

Figure 2.10: Event study: annual labour earnings

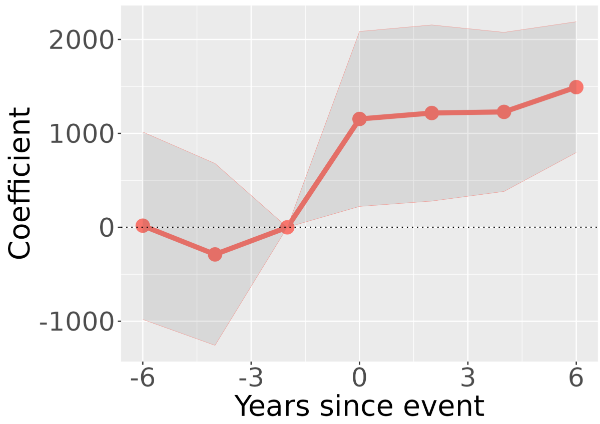

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Finally, Figure 2.11 plots the estimates where the LHS variable is spouse's labour force participation. Again, spouse's labour force participation seems to increase after the mistake about the SPA is realised, though there are some pre-trends here as well.

Figure 2.11: Event study: spouse's labour force participation

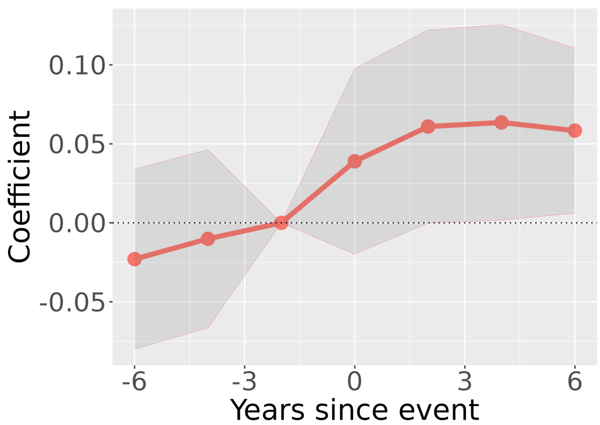

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

However, there is no corresponding uptick for other variables of interest. Figure 2.12 below plots the event studies where the outcome variables are log annual food spending, a dummy for being a homeowner, a dummy for having over £10k in financial wealth and a dummy for financially struggling[12].

In no case is there a striking uptick or downtick in the outcome variable around the moment of realisation. As such, to the extent that we are able to interpret these event studies as measuring responses to realising one's pension will arrive later than expected,

[12] ELSA respondents are asked "Looking at this card, please say how often you find you have too little money to spend on what you feel your needs are?", where the options are "Never", "Rarely", "Sometimes", "Often" or "Most of the time". I classify as financially struggling anyone who answers "Sometimes", "Often" or "Most of the time".

Figure 2.12: Event studies: other outcome variables

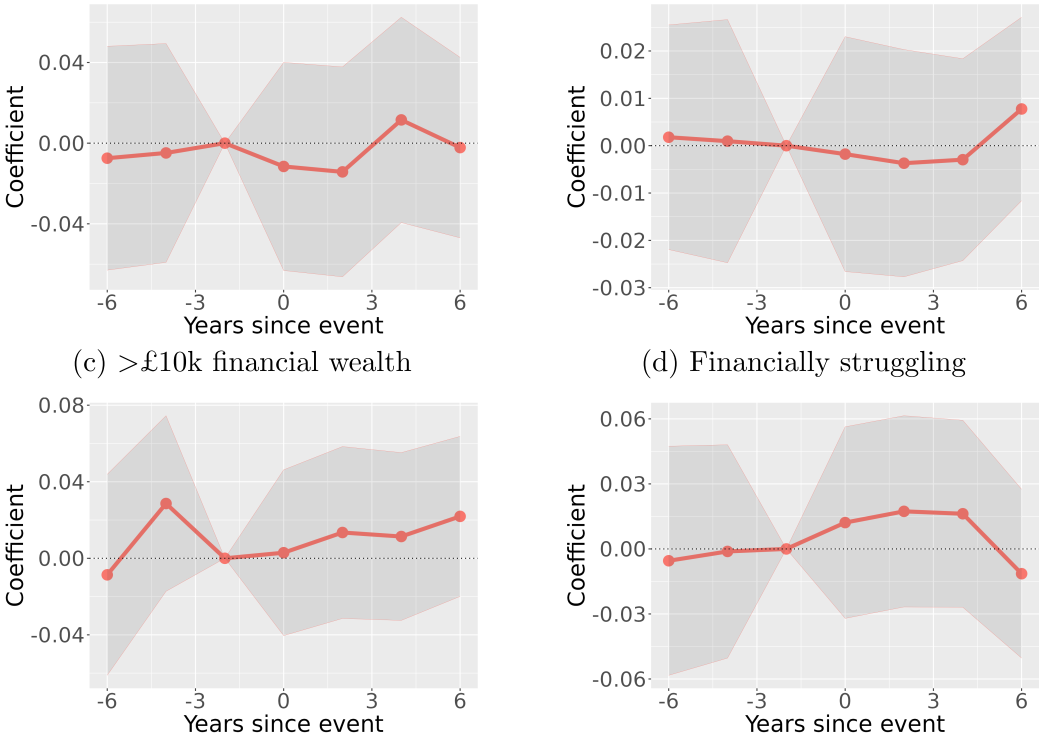

(a) Log annual food spending

(b) Own home

(c) >£10k financial wealth

(d) Financially struggling

Notes: standard errors clustered at the individual level.

most of the response is loaded on changes in labour supply rather than a drawing down of assets or a reduction in consumption, and that there is no evidence of a deterioration in households' financial situation.

However, as outlined above, the significant problem with interpreting these results causally is that the arrival of information is in part a result of agents' endogenous decisions to do their own research. For this reason, it is important to model the incentives agents face to do their own research.

2.2.3.1 Robustness

In Appendix 2.B.4 I present some supplementary event study graphs using different outcome variables and using different controls and assumptions about trends. In each case of changing controls or assumptions about trends the results are qualitatively similar. I also present some graphs showing proportions of agents with correct beliefs over time and over the age distribution.

2.3 Model

2.3.1 Model overview

The decision-maker in the model is a household, either a single woman or a married woman-man couple. The household's age is set to be the same as the age of the woman. The household starts the model at age 50 and dies at age 85. If there is a husband in the household, his age is assumed to be 2 years greater than that of the woman, matching median age differences in the data. Time is discrete.

In every period, the household is either attentive or inattentive to the true SPA. If the household is attentive, then the model collapses to a standard life-cycle model with a unitary household where each period the household makes a discrete labour supply choice and a continuous consumption choice, saves in a risk-free asset, and receives state pension income from the government from the true SPA onwards, which they fully anticipate.

If the household is inattentive then they face an extra choice at the beginning of each period: they choose their level of attention $m_t \in \{0, 1\}$, which can be intuitively thought of as the household "doing their research" once and for all. The benefit of doing research is that the household resolves their uncertainty about their true SPA and becomes attentive, enabling better planning for retirement; the cost of doing research is that the household must pay a cost of attention which depends on their state variables. Once a household does their research, they stay attentive forever, so the model again collapses into a more standard life-cycle model with consumption and labour choices.

2.3.2 Beliefs

A key element of the model is a household's[13] beliefs about their true SPA. Let the belief of an inattentive agent be given by g^0, a vector of probabilities associated with each possible value of the true SPA. To keep the problem tractable, I impose a common prior across all inattentive agents, and I assume that the agent believes that the SPA can only take on integer values between 60 and 66 inclusive. For instance, g^0 might be represented by the probability mass function in Figure 2.13.

In this case, the inattentive agent is confident that the true SPA is 60 but allows some possibility that it will take on higher values.

By contrast, the attentive agent is certain of what the true SPA is. Their belief is given by $\bar{g}$, a degenerate probability mass function which assigns a probability of 1 to the actual true SPA being the SPA.

If the agent were to do their research and become attentive, their beliefs would switch

[13]Given that, at least in earlier waves of ELSA, only women were asked about their SPA, I do not differentiate between the beliefs of the woman or a man in a couple household, and will refer to agents'/households' beliefs interchangeably.

Figure 2.13: Example g^0

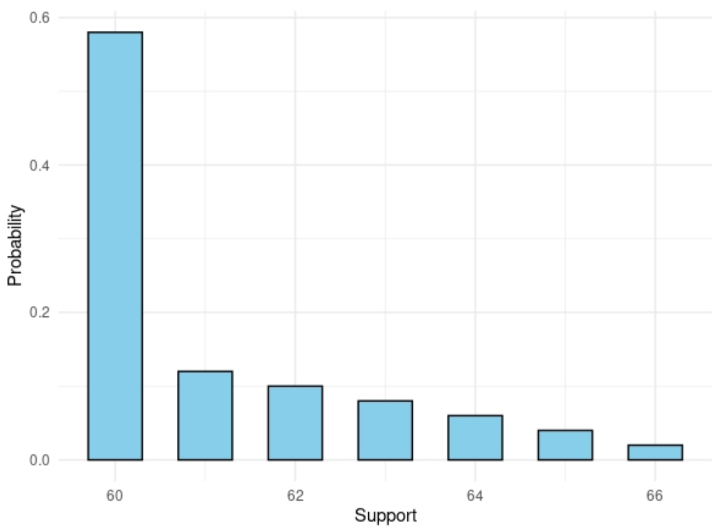

from g^0 to $\bar{g}$. In the inattentive state, they therefore have beliefs about what their beliefs will be should they pay attention, given by g^0. For instance an inattentive agent might think that there is an 80% chance that if they pay attention they will subsequently be certain that the true SPA is 60. Thus the prior belief g^0 will play a key role in determining the inattentive agent's incentives to do their research.

2.3.3 Resources

Households earn wages in the labour market and save in a risk-free asset. Let z_t be a collection of state variables and let c_t and L_t be the household's consumption and labour market choices, outlined in Section 2.3.4 below. The household budget constraint is:

$$a_{t+1} = (a_t + h_l(L_t, z_t) + h_{nl}(z_t) - c_t)(1 + r) \tag{2.2}$$

where a_t is wealth, $h_l(L_t, z_t)$ is labour market income, $h_{nl}(z_t)$ is non-labour market income and r is the real interest rate.

If the woman in the household does not work in the labour market, she earns nothing; otherwise, she earns a wage $w_f(z_t)$. Similarly, if there is a man in the household and he works then he earns a wage $w_m(z_t)$. Household labour income $h_l(L_t, z_t)$ is the sum of wages across the household members. In any given period, wages will depend on household members' demographics as well as temporary productivity shocks, independent across household members.

As for non-labour income, household members receive their state pensions from their respective SPAs. Some household members may also receive private pension income from that pension's normal retirement age. Households also receive a residual amount of non-labour income that does not come from pensions, which captures capital income, government transfers and other sources of income. For full parameterization of both the labour market and non-labour-market income functions see Section 2.3.5 below.

2.3.4 Choices

2.3.4.1 Consumption choice

Working backwards within a given period, the final choice that a household makes is consumption, c_t. Conditional on state variables z_t, labour choice L_t, and current belief about the SPA g_t, they choose c_t to maximise:

$$v_t^c(c_t|L_t, z_t, g_t) = u(c_t, z_t) + \beta \mathbb{E}_{g_t}\left[V_{t+1}^m(z_{t+1})|L_t, c_t, z_t\right] \tag{2.3}$$

where $u(c_t, z_t)$ is the instantaneous utility function for consumption and $V_t^m(\cdot)$ is the value function from the attention choice, defined below. Thus, as is standard, households choose consumption to maximise the expected present discounted value of lifetime utility, where the expectations are taken relative to their current beliefs g_t (hence the g_t subscript on the expectations term).

2.3.4.2 Labour choice

Prior to the consumption choice the household makes a discrete labour choice. Conditional on state variables z_t and current belief g_t, the agent chooses one of four options for labour supply: {Neither wife nor husband works; Wife does not work but husband does; Wife works but husband does not; Both work}. Denote these four possible choices as $L_t = \{1, 2, 3, 4\}$. For single women, only the first and third of these are available as choices.

Let $V_t^c(L_t, z_t, g_t) = \max(v_t^c(c_t|L_t, z_t, g_t))$, i.e. the value function from the optimal consumption choice. The utility associated with labour choice L_t is given by:

$$v_t^L(L_t|z_t, g_t) = V_t^c(L_t, z_t, g_t) - \omega(L_t, z_t) + \epsilon_{Lt} \tag{2.4}$$

Thus, the utility from the labour choice is the value function from the consumption choice conditional on that labour choice, less cost of labour $\omega(L_t, z_t)$, plus an iid Type 1 Extreme Value preference shock with scale parameter σ_ϵ[14]. The cost of supplying labour is allowed to depend on state variables z_t because, for instance, it might be more costly to work in t if one did not work in $t-1$, because of start-up costs.

[14] I assume that the labour preference shocks are realised only after the agent makes their attention choice, prior to the labour choice. Thus, when the agent is making their attention choice, they only know the probabilities of their future labour supply.

2.3.4.3 Attention choice

Prior to the labour choice, if the agent is inattentive, they make an attention decision $m_t \in \{0, 1\}$. As set out above, this is tantamount to the agent deciding whether to do their research "once and for all" on the issue of the SPA and resolving any uncertainty; in other words, being attentive is an absorbing state.

Let $V_t^L(z_t, g_t) = \mathbb{E}_{g_t}[\max(v_t^L(L_t|z_t, g_t)]$, i.e. the value function from the optimal labour choice. The utility of setting $m_t = 1$ is given by:

$$v_t^m(1|z_t) = \mathbb{E}_{g^0}[V_t^L(z_t, \bar{g})] - \kappa(z_t) + \xi_{1t} \tag{2.5}$$

and of setting $m_t = 0$

$$v_t^m(0|z_t) = V_t^L(z_t, g^0) + \xi_{0t} \tag{2.6}$$

where $\kappa(z_t)$ is the cognitive cost of paying attention and ξ_{it} is another iid Type 1 Extreme Value preference shock, independent of the labour preference shock, with spread parameter σ_ξ.

To close Equation 2.3 above, we need to specify $V_t^m(z_t)$. For inattentive agents, this will be $V_t^m(z_t) = \mathbb{E}_{g^0}[\max(v_t^m(m_t|z_t))]$, i.e. the value function from making the optimal attention choice. For attentive agents, as they have no attention choice to make, we have that $V_t^m(z_t) = V_t^L(z_t, \bar{g})$. In the terminal period, there is no future utility, so $V_{t+1}^m(z_{t+1}) = 0$.

Equation 2.5 suggests that it is the expectation of the value function from the labour choice which is key for assessing agents' incentives to pay attention. This is because an inattentive agent is uncertain over what the degenerate distribution $\bar{g}$ is, with different values for $\bar{g}$ resulting in different values for the value function. Thus, the fact that inattentive agents evaluate this expectation according to their original belief distribution g^0 is crucial.

I assume that if the inattentive agent reaches the age of 60, and the true SPA for this agent is indeed 60, then their uncertainty is resolved costlessly: they become attentive without incurring attention cost $\kappa(z_t)$. Otherwise, they are forced to pay cost $\kappa(z_t)$ and become attentive. This captures the fact that if an inattentive agent reached their expected SPA and ended up receiving their pension, their inattention would be costlessly resolved, whereas the inattentive agent who did not receive their pension at their expected SPA would be forced into doing their own research.

2.3.4.4 Discussion

Another way of stating the problem of whether to pay attention for the inattentive household is that the inattentive household chooses to do their own research iff:

$$\underbrace{\mathbb{E}_{g^0}[V_t^L(z_t,\bar{g})] - \kappa(z_t) + \xi_{1t}}_{\text{Value of resolving uncertainty}} > \underbrace{V_t^L(z_t,g^0) + \xi_{0t}}_{\text{Value of delaying resolution of uncertainty}} \tag{2.7}$$

This condition is intuitively more likely to hold if:

- The cost of attention $\kappa(z_t)$ is lower. For instance, people in couples plausibly have a lower cost of attention because there is an extra person in the household who might come across information on SPA changes, so other things being equal we might expect these people to pay more attention.

- The "stakes" of the problem are higher, i.e. optimal behaviour is highly sensitive to when the true SPA is, and/or utility is very sensitive to optimal behaviour. For instance, a woman with a private pension has an extra source of income in the household, so is better insured against future shocks to household income: the woman's state pension income is a smaller fraction of total household income in retirement. This means that the stakes of the issue are lower for such households, and so other things being equal we might expect these people to pay less attention.

- There is more uncertainty in the original prior belief g^0. If inattentive households are originally less certain that the true SPA is 60 then they will give more weight to the "bad" outcome where they have to wait several years for their pension, under which carrying on the same path of consumption and labour as before might be highly suboptimal. They then will have more of an incentive to do their research and find out what the true SPA is.

- Most simply, there is a large positive net preference shock for paying attention $\xi_{1t} - \xi_{0t}$. These shocks could be interpreted as genuine iid variation in preferences for doing research. However, they could also be interpreted as exogenously occurring opportunities to receive information - like an article appearing in a local newspaper - which e.g. couples might be more likely to take on board because there is an extra set of eyes in the household[15].

2.3.5 Functional forms and parameterisation

2.3.5.1 Belief parameters

The key object for agents' beliefs is the original belief distribution g^0, which is the probabilities that inattentive agents attach to the true SPA taking on certain values between 60 and 66. From the survey data, we only know what inattentive agents consider their

[15]Due to this ambiguity over the interpretation of shocks to the cost of paying attention, the welfare analysis in Section 2.5 will largely abstract from utility from attention and focus instead on utility from consumption and labour.

true SPA to be, but we do not know the certainty with which they hold this belief, or how likely they are to consider themselves to be wrong. To capture this potential uncertainty simply within the model, I impose, for $k \in \{60, ..., 65\}$, that:

$$P(SPA = k + 1) = \lambda P(SPA = k) \tag{2.8}$$

where $P(SPA = k)$ is the agent's subjective probability under g^0 that the true SPA for them is at age k. This, combined with the assumption that $\sum_{k=60}^{66} P(SPA = k) = 1$, means that the entire distribution g^0 is determined by a single parameter λ, which captures the degree of uncertainty that agents have over the true SPA when inattentive. As $\lambda \to 0$, inattentive agents are certain that the true SPA is 60, and as $\lambda \to 1$, inattentive agents attach equal probability to the SPA being any age from 60 to 66.

2.3.5.2 Utility of consumption

I assume that utility from consumption takes the CRRA form:

$$u(c_t, z_t) = \frac{\left(\frac{c_t}{\chi(z_t)}\right)^{1-\theta}}{1-\theta} \tag{2.9}$$

where $\chi(z_t)$ is the equivalence scale for consumption, which I set as 1 for singles and $\sqrt{2}$ for couples.

2.3.5.3 Disutility of working

The disutility of working is captured in parameter $\omega(L_t, z_t)$.

For each individual household member $i \in \{m(\text{ale}), f(\text{emale})\}$, let L_t^i be the household member's labour supply at t. Let $\omega_i(L_t^i, z_t)$ be that individual's contribution to total household disutility of working $\omega(L_t, z_t)$. Then, I impose that $\omega_i(L_t^i, z_t)$ takes the form:

$$\omega_i(L_t^i, z_t) = \begin{cases} \omega_0 + \omega_{age} \times age_{it} + \omega_{age2} \times age_{it}^2 + \omega_{prevwork} \times L_{i,t-1} & \text{if } L_t^i = 1 \\ 0 & \text{otherwise} \end{cases} \tag{2.10}$$

As such, the disutility of working depends on a quadratic in age and on whether the household member was working the previous period.

For single households, we have simply that $\omega(L_t, z_t) = \omega_i(L_t^i, z_t)$[16]. For couple households, I impose that:

[16]Where $L_t = 1$ if $L_t^i = 0$ and $L_t = 3$ is $L_t^i = 1$, as L_t corresponds to the set {Neither wife nor husband works; Wife does not work but husband does; Wife works but husband does not; Both work} whereas L_t^i corresponds to the set {Household member i does not work; Household member i works} .

$$\omega(L_t, z_t) = \begin{cases} \omega_f(L_t^f, z_t) + \omega_m(L_t^m, z_t) + \omega_{joint} & \text{if } L_t^f = 0 \ \& \ L_t^m = 0 \\ \omega_f(L_t^f, z_t) + \omega_m(L_t^m, z_t) & \text{otherwise} \end{cases} \tag{2.11}$$

In other words, for couple households, the overall disutility of labour is the sum of individual members' disutilities plus an extra interaction term for the case where neither household member is working, which would allow the model to capture potential costs or (if negative) benefits of both household members retiring at the same time, as discussed in e.g. Lalive and Parrotta (2017).

2.3.5.4 Cost of paying attention

The cognitive cost of paying attention is captured in $\kappa(z_t)$. I parameterise it as follows:

$$\kappa(z_t) = \kappa_0 + \kappa_{HS} \times highschool + \kappa_{col} \times college + \kappa_{couple} \times couple + \kappa_{prevwork} \times L_{t-1}^f \tag{2.12}$$

Thus, the cost of paying attention is a linear function of education level, couple status, and whether the woman in the household worked the previous period.

2.3.5.5 Wages and non-labour income

In this subsection I discuss the parameterization of the wage and non-labour income processes.

I assume the wage on offer for household member i every period is:

$$\begin{aligned} w_i(z_t) = \exp(&\gamma_0 + \gamma_{age} \times age + \gamma_{age2} \times age^2 + \gamma_{HS} \times highschool_i \\ &+ \gamma_{col} \times college_i + \gamma_{male} \times male_i + \gamma_{prevwork} \times L_{t-1}^i + \eta_{it}) \end{aligned} \tag{2.13}$$

where

$$\eta_{it} = \rho_i \eta_{i,t-1} + \nu_{it} \tag{2.14}$$

is a persistent productivity shock for agent i, modelled as an AR(1)[17]. The shock component ν_{it} is iid over time and across household members and has distribution $N(0, \sigma_{\nu i}^2)$. Note that ρ_i and $\sigma_{\nu i}^2$ have i subscripts because the persistence and variance of productivity shocks can differ by gender to allow for different wage processes for men and women.

[17]When it comes to solving the model, in order to model the AR(1) part of the process, I use the Rouwenhorst method (Kopecky and Suen 2010) to discretize the AR(1) and convert it into a Markov process.

Households have three potential sources of non-labour income: state pension, private pension, and other income, with this latter category capturing other government transfers and capital income. Households are assumed to have private pensions if and only if they are ever observed receiving or contributing to private pensions in the data.

I assume state pension income is a flat amount, which is received from the household member's true SPA onwards: generally 65 for men and some age between 60 and 65 for women, depending on cohort. The flat amount is taken from mean state pension income in the data.

For those with private pensions, private pension income is also a flat amount, and is received from age 60 for women and age 65 for men. I allow private pension income to differ by the education status of the household member and the couple status of the household, and use mean private pension income received by these demographic groups in the data.

Finally, other income is received every period at a flat rate, which I allow to differ by household couple status and by the education levels of wife (and husband, if relevant). The flat rate is then given by the mean non-labour and non-pension income received in the data by households with the relevant demographics.

Then, household non-labour income is simply given as the sum of these three income types across household members.

2.3.5.6 Model solution

The time-varying state variables of the model are the household's age, liquid wealth, beliefs at the start of the period, productivity shock of woman (and husband, if any) and labour supply previous period of woman (and husband, if any). The non-time-varying state variables are then the couple status of the household, education of the woman (and husband, if any)[18] and whether or not the woman (and/or husband) have a private pension.

The model is solved via backwards induction. First, for each candidate SPA between 60 and 66, I solve the model for every combination of state variables, imposing that agents are certain that the true SPA is the candidate SPA. Then, I use the value functions from this process as inputs for solving the problem of the inattentive agent, constructing the equivalents of Equations 2.5 and 2.6 to determine the incentives for the inattentive agent to pay attention when they are below the SPA.

[18] There are three separate categories for education, namely leaving high school at or below 16, having a high school education and having a college education.

2.4 Estimation

Estimation takes place using the data from Wave 3 (2006) onwards for the 1945-56 birth cohort ELSA sample of women and their husbands (if any). I select only those women who were present in Wave 3 and who have information on their work status, income and (if below SPA) pension beliefs for each subsequent wave that they answer. As the model does not allow for households' couple status to change, I drop any households who have their couple status change over the period that they are observed in the data (10.7% of all households). Appendix 2.B.1 presents more detail on the sample used for estimation. I impose that agents leave the model in the same period that they leave the data in order to ensure the same attrition affects both data moments and simulated moments.

I use as the relevant wealth variable the sum of financial wealth for the household, meaning I exclude housing wealth. Housing wealth is a very large component of total wealth of English households in retirement: in my sample, the mean (median) housing wealth per household at age 60 is £276k (£199k), compared to £80k (£25k) for financial wealth. However, in the UK, downsizing or otherwise accessing this housing wealth is rare (Banks, Blundell, et al. 2012; Blundell, Crawford, et al. 2016), meaning that practically the wealth that households will have access to in order to smooth any income shocks will be their more liquid financial wealth[19]. Moreover, modelling housing and financial wealth separately would severely complicate the model. For this reason, I focus only on financial wealth. However, in Appendix 2.B.7, I re-estimate the model and the main welfare analysis using total wealth, i.e. financial wealth plus housing wealth, rather than only financial wealth, as the wealth variable. Unsurprisingly, the results are that working is associated with a smaller utility penalty (in order for people to work realistic amounts despite their greater wealth) and the welfare costs of inattention are much smaller (because households have so much more wealth to protect against falls in consumption).

2.4.1 Parameters from outside the model

Some parameters are set outside of the model for convenience. I assume that the real interest rate r is constant at 2% and that the discount factor is $\beta = 0.975$. I set the CRRA parameter θ equal to 2.

As for the state pension ages for each cohort, to make the problem tractable I divide the data into 6 2-year birth cohorts, from those born 1945-46 (Cohort 0) to those born 1955-56 (Cohort 5). I set each agent's initial SPA as the median SPA for their cohort according to the regime in place in 2006, the first wave of the data, though I allow for

[19] Indeed, the event studies of Figure 2.12 do not point to people selling their houses or otherwise converting their housing wealth into financial wealth at the point of realising their erroneous pension beliefs.

each cohort's SPA to change following the 2011 Pensions Act outlined in Figure 2.2 and surrounding discussion[20].

As set out in Section 2.3.5.5, state pension, private pension and other income sources are set at their mean levels in the data by demographic group.

2.4.2 Parameters estimated inside the model

The parameters to estimate inside the model are summarised in Table 2.1 below. There are 24 parameters to estimate in total.

Estimation takes place via the Method of Simulated Moments. I choose the parameter vector $\hat{\theta}$ to minimize the objective function:

$$(\bar{\mu} - \mu(\theta))' W (\bar{\mu} - \mu(\theta)) \tag{2.15}$$

where $\bar{\mu}$ is a vector of moments from the data, $\mu(\theta)$ is vector of corresponding moments generated by the model under parameter vector θ, and W is a weighting matrix.

I target 5 main different types of moments:

- the mean labour earnings of women;
- the mean labour earnings of men;
- the proportion of women working;
- the proportion of men working;
- the proportion of women with correct beliefs about their SPA[21].

I target these moments for every wave from 2006 to 2018, separately for each of the following subgroups:

- each birth cohort
- each education level (of both wife and husband, if any)
- each private pension status (of both wife and husband, if any)
- each labour force status in the previous period (for both wife and husband, if any)
- single vs. married women.

[20] Specifically, I allow for each cohort's SPA to unexpectedly change in 2011, i.e. the 6th period of the model after starting in 2006. I assume that all agents who were aware of the pre-2011 pension system are aware of the 2011 reforms and update their SPA beliefs accordingly, whereas all those who were unaware before 2011 continue being unaware afterwards, with the same prior g^0, until they correct their beliefs.

[21] I only consider women younger than 60 for this proportion because those 60 or above mechanically will have correct beliefs, because I assume that women's uncertainty is automatically resolved when they reach the old SPA.

Table 2.1: Parameter Descriptions

Parameter	Description
Work penalty parameters	
ω_0	Work penalty: constant
ω_{age}	Work penalty: linear age trend
ω_{age2}	Work penalty: quadratic age trend
$\omega_{prevwork}$	Work penalty: worked previous period
ω_{joint}	Work penalty: taking leisure at same time as spouse
Wage parameters	
γ_0	Log wage: constant
γ_{age}	Log wage: linear age trend
γ_{age2}	Log wage: quadratic age trend
$\gamma_{prevwork}$	Log wage: worked previous period
γ_{male}	Log wage: male
γ_{HS}	Log wage: high school education
γ_{col}	Log wage: college education
Attention cost parameters	
κ_0	Attention penalty: constant
$\kappa_{prevwork}$	Attention penalty: worked previous period
κ_{HS}	Attention penalty: high school education
κ_{col}	Attention penalty: college education
κ_{couple}	Attention penalty: in couple
Other parameters	
$\sigma^2_{\nu w}$	Variance of iid component of prod. shock, women
ρ_w	Persistence of prod. shock, women
$\sigma^2_{\nu m}$	Variance of iid component of prod. shock, men
ρ_m	Persistence of prod. shock, men
λ	Decay factor for beliefs
σ_ϵ	Spread of preference shock for labour
σ_ξ	Spread of preference shock for attention

I drop any moment from the estimation where there are fewer than 20 observations in the data to construct the relevant mean. I am left with 601 base moments. To these I add six supplementary moments to match:

1) the coefficient on lagged log labour earnings in a regression of women's log labour earnings on lagged log labour earnings, individual fixed effects and age, conditional on working in both periods;

2) the variance of the residuals from the above regression;

3) the equivalent of (1) for men's lagged log labour earnings;

4) the equivalent of (2) for men's lagged log labour earnings;

5) the coefficient on a dummy for having correct beliefs about one's SPA in a regression of women's labour force participation on the correct beliefs dummy, individual fixed effects and age;

6) the coefficient on a dummy for husband's labour force participation in a regression of a wife's labour force participation on the dummy for their husband's labour force participation, individual fixed effects and age (for couples only).

I discuss the reasons for including these moments in particular in Section 2.4.3 below.

The baseline weighting matrix I use is a diagonal matrix where the off-diagonal elements are 0 and the diagonal elements are the inverse of estimates of the variance of the corresponding data moment. This allows more precisely estimated data moments to receive more weight. However, to avoid the six supplementary moments from being swamped by sheer force of numbers by the 601 baseline moments, I reweight so that the supplementary moments receive 1/6 of the total weight and the baseline moments receive 5/6, reflecting the fact that the baseline moments represent 5 basic moment types, with each moment type receiving approximately the same weight as the supplementary moments taken together.

2.4.3 Identification

In this section I provide informal arguments for how the parameters to be estimated are identified from the data.

The parameters governing mean wages by age and by demographics are identified by differences in wages across ages and between different demographic groups. Then, the variance and autocorrelation of shocks to wages for men and women are identified by the first four supplementary moments described in the previous section, i.e. the variance of the residuals and the coefficient on lagged wages in wage regressions for men and women separately.

The parameters governing the utility penalties of working are identified by the proportions of agents working in the data. The fact that we observe people's wages when working mean that we can separate out people leaving work because the work penalty is increasing and because their wage is declining. Moreover, we can identify the parameter governing the utility bonus husbands and wives receive from not working at the same time using the sixth supplementary moment of the previous section, i.e. the coefficient on husband's labour force participation in a regression of wife's labour force participation on husband's participation, individual fixed effects and age.

The parameters governing the utility penalties of paying attention are identified by the proportions of agents in different demographic groups who have correct beliefs in the data. For instance, if married women are more likely to have correct beliefs even if the stakes of the paying attention are less for them (because their state pension is a smaller share of household income in retirement), then this would suggest that married women face a lower cost of paying attention than single women.

All utility cost and bonus parameters, including the two discrete choice shock spread parameters, σ_ϵ and σ_ξ, are scaled relative to the coefficient multiplying the CRRA utility function over consumption, which is normalised to 1. The larger is σ_ϵ (the labour preference shock), the smaller the gap in labour force participation rates between those with a large incentive to work and those with a smaller incentive to work, because more of the labour supply decision is explained by unobserved preference shocks. For instance, other things being equal, women without a private pension have a larger incentive to work than women with a private pension, because women with a private pension have an alternative income stream to rely on for retirement. Yet, women with a private pension have no difference in the utility penalty of working relative to women without. Thus, if women without a private pension are much more likely to work than women with a private pension, this suggests σ_ϵ is small, whereas if the gap is less pronounced, σ_ϵ would be larger. A similar argument applies to comparisons between more and less educated women, and women with and without a husband.

Similarly, to identify σ_ξ, we can match gaps in mean attention rates between different groups who are assumed to have the same utility penalties of paying attention, but have different incentives to pay attention for other reasons. In particular, it is assumed that women with a private pension will find it no more difficult to pay attention, yet women with a private pension have less of an incentive to pay attention to the SPA, so if the gap in attention rates between women with and without a private pension is very large, that would suggest σ_ξ is small.

Finally, the belief decay parameter governs how confident inattentive people are that the true SPA is 60. The more certain someone is that the true SPA is 60, then the less likely they will be to pay attention, and the bigger the correction in their behaviour after they realise. In contrast, the attention utility parameters will affect the cost of paying attention (hence the probability of someone paying attention) but will not affect how big the adjustment in behaviour will be after someone starts paying attention. Therefore, the bigger the adjustment in behaviour is following a realisation, the more certain an agent was in their beliefs beforehand, other things being equal.

2.4.4 Results

The results of the estimation are presented in Table 2.2 below.

The work penalty parameters capture variation between agents in the utility loss they suffer if they work. Notably, the penalty associated with working is much less if the agent

Table 2.2: Estimation Results

Parameter	Estimate
Work penalty parameters	
ω_0 - Constant	2.300
	(0.114)
ω_{age} - Linear age trend	-0.015
	(0.001)
ω_{age2} - Quadratic age trend	0.002
	(0.000)
$\omega_{prevwork}$ - Worked last period	-2.179
	(0.131)
ω_{joint} - Joint leisure	-0.202
	(0.011)
(Log) Wage parameters	
γ_0 - Constant	-0.405
	(0.022)
γ_{HS} - High school education	0.269
	(0.034)
γ_{col} - College education	0.766
	(0.033)
γ_{age} - Linear age trend	-0.024
	(0.002)
γ_{age2} - Quadratic age trend	-0.001
	(0.000)
$\gamma_{prevwork}$ - Worked last period	0.183
	(0.002)
γ_{male} - Male	0.232
	(0.038)
Attention penalty parameters	
κ_0 - Constant	1.784
	(0.038)
κ_{HS} - High school education	-0.018
	(0.001)
κ_{col} - College education	-0.273
	(0.011)
$\kappa_{prevwork}$ - Worked last period	-0.103
	(0.015)
κ_{couple} - Couple	-0.307
	(0.021)
Other parameters	
$\sigma^2_{\nu w}$ - Variance of iid component of prod. shock, women	0.557
	(0.032)
ρ_w - Persistence of prod. shock, women	0.414
	(0.032)
$\sigma^2_{\nu m}$ - Variance of iid component of prod. shock, men	0.741
	(0.036)
ρ_m - Persistence of prod. shock, men	0.545
	(0.017)
λ - Belief decay parameter	0.347
	(0.025)
σ_ϵ - Type 1 EV spread parameter, labour	0.549
	(0.026)
σ_ξ - Type 1 EV spread parameter, attention	0.467
	(0.020)

Notes: estimation via MSM. See Section 2.4.2 for discussion of moments and weighting matrix used. Standard errors calculated from 20 bootstrap replications. For the calculation of age trends, note that ages were normalised such that age 50 in the data was age 0 in the model.

worked in the previous period, reflecting the fixed costs involved in starting work again having dropped out of the labour force.

To interpret the size of these coefficients, one can consider that if a given 50 year old had a 50% chance of working in a given period, then an agent who was entirely equivalent other than facing a 70 year old's cost of working today would instead work with only 28.7% probability. This comes from the fact that if Agent A has a 50% chance of working, then $v_A = 0$, where v_A is the value of working relative to not working, and thus if for Agent B $v_B = v_A + 0.015 \times (70 - 50) - 0.002 \times (70 - 50)^2$ and if $\sigma = 0.549$ is the spread parameter for the relevant Type 1 EV shock, then the probability of B working is $P_B = \frac{\exp(v_B/\sigma)}{1+\exp(v_B/\sigma)} = 0.287$. Similarly, if Agent A was identical to Agent B apart from the fact that Agent A had worked the previous period, and Agent A had a 50% chance of working today, then Agent B would have only a 1.8% probability of working today.

The ω_{joint} parameter captures the extent to which couples receive a negative utility penalty, hence a utility boost, if they are not working at the same time; in this case, if Agent A has a 50% chance of working when their spouse is working, they will have a 40.9% chance of working if their spouse is not working, other things being equal.

The wage parameters capture the difference in log potential wages, or (approximately) the percentage difference in potential wages, between different groups. Those with a high school education receive approximately 26.9% more in wages than those who left school at 16, whereas those with a college education receive 76.6% more. Wages decline quadratically with age, at a rate of approximately 2.4% per year at age 50 and 4.4% per year at age 60. Again, those who worked in the previous period receive 18.3% higher wages, and men receive 23.2% higher wages than women.

The attention penalty parameters capture the utility loss of paying attention and "doing one's research" on the reforms to the state pension age. In other words, these parameters capture the cost of attention for different groups of women[22]. High-school educated women face similar costs of attention to women who left school at 16, but women with a college education have notably lower costs. Using the same approach as before, if a woman who left school at 16 had a 50% chance of becoming attentive in a given period, then a woman who was identical apart from the fact that her attention costs today was that of a high school educated woman would have a 51.0% chance of paying attention, and a woman who had an attention cost today of a college-educated woman would have a 64.2% chance of paying attention. Those who worked in the previous period have slightly lower costs of paying attention, and those in couples have significantly lower costs.

The variance and persistence of the productivity shock are estimated to match the variance and persistence in the wage data for men and women separately. Men's wages

[22]The attention parameters are relevant to women only because only women are inattentive about their SPA, because men's SPA is not changing.

show higher variance and persistence of productivity shocks. The belief decay parameter captures how sure inattentive agents are that their true SPA is 60. The value of 0.347 means that for an inattentive agent, $Pr(SPA = k + 1) = 0.347Pr(SPA = k)$, for k between 60 and 65. Finally, the Type 1 EV shock spread parameters capture how much *iid* variation in preferences over time there is which is not captured by the model: there is more unobserved iid heterogeneity in the labour choice than in the attention choice.

2.4.5 Model fit

Here, I present the targeted moments pooling over all cohorts to give a summary indication of model fit. In Appendix 2.B.5 I show fit disaggregated across different cohorts[23].

Figure 2.14 shows women's earnings, in the data and in the model. The model is able to broadly match the level and (downward) trajectory of women's earnings.

Figure 2.14: Model fit: women's earnings

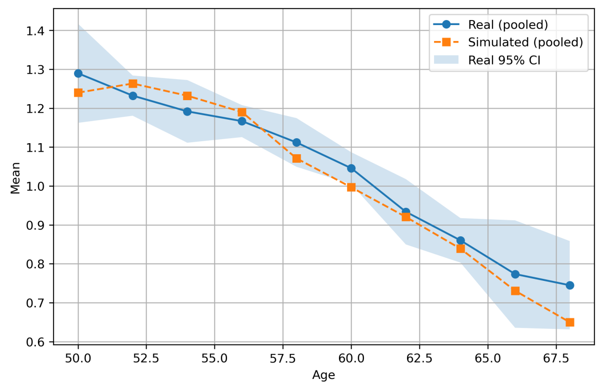

Figure 2.15 shows the proportion of women in work in the data and in the model. Again, the model captures well the broad trajectory of the proportion of women working, with particularly steep declines in working rates around age 60 in the model and in the data.

[23]Note that as well as matching the targeted moments by cohort, I also match them by education status, private pension status, previous labour force status and couple status, as set out in Section 2.4.2 For reasons of space, I focus only on the targeted moments by cohort as they provide the most useful summary measure of model fit.

Figure 2.15: Model fit: proportion of women working

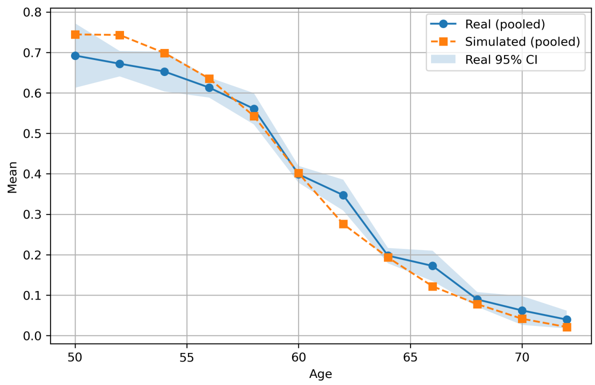

Figure 2.16 shows men's earnings in the data and in the model. Here, the model broadly matches the (slightly noisier) mean earnings data by age for men.

Figure 2.16: Model fit: men's earnings

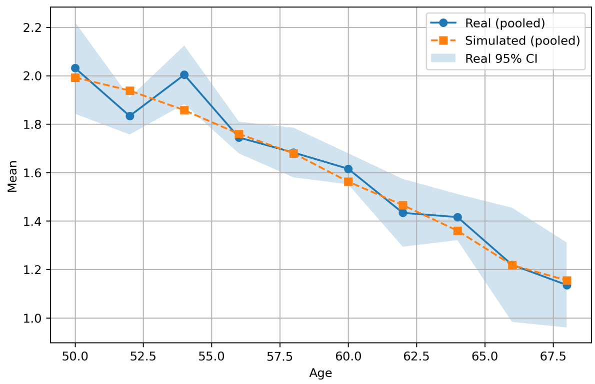

Figure 2.17 shows the proportion of men in work in the data and in the model.

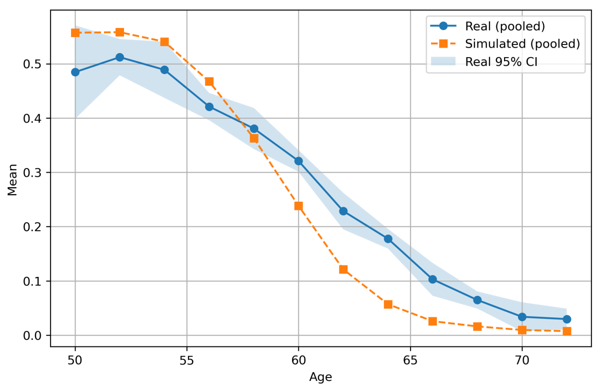

Figure 2.17: Model fit: proportion of men working

Here are the most significant failures of model fit - the model underestimates the proportion of men working in their 60s, and in general predicts a steeper decline in working rates than is observed in the data. This is driven by the fact that, as the disaggregated figures by cohort in Appendix 2.B.5 suggest, women's labour force participation in the data exhibits notable cohort effects which are not observed to the same degree in men's labour force participation, and which cannot entirely be explained by changing SPAs. In order to match these cohort effects for women, who make up more of the sample relative to men[24], the model decreases the cost of working in one's 50s (i.e. for younger generations at the time of observation) and increases the cost of working in one's 60s (i.e. for older generations). As there are no corresponding cohort effects for men, the result is that the model overestimates men's working rates in their 50s and underestimates men's working rates in their 60s.

Figure 2.18 shows the proportion of households who are aware of the pension reform in the data and the model. Again, the model broadly matches the level and trajectory of attention paid by households in the data.

Figure 2.18: Model fit: attention

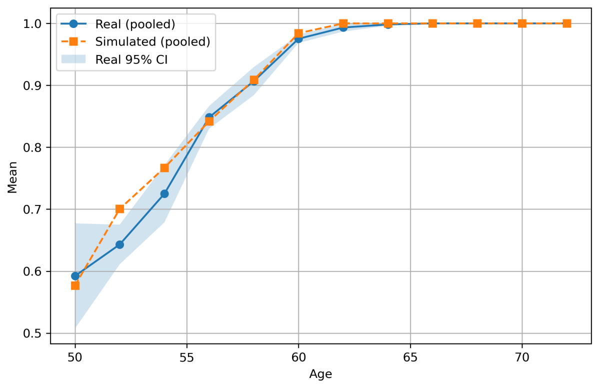

Figure 2.19 shows the supplementary moments:

1) the coefficient on lagged log labour earnings (2 years ago)[25] in a regression of women's log labour earnings on lagged log labour earnings, individual fixed effects and age, conditional on working in both periods;

2) the variance of the residuals from the above regression;

3) the equivalent of (1) for men's lagged log labour earnings;

4) the equivalent of (2) for men's lagged log labour earnings;

5) the coefficient on a dummy for having correct beliefs about one's SPA in a regression of women's labour force participation on the correct beliefs dummy, individual fixed effects and age;

[24]This is because households in the model and data are either couple households or single women.

[25]As ELSA waves are biennial but each period in the model is a year, the regression using model-generated data needs to use lagged log wage from 2 periods ago to be comparable with the original ELSA data.

6) the coefficient on a dummy for husband's labour force participation in a regression of a wife's labour force participation on the dummy for their husband's labour force participation, individual fixed effects and age (for couples only).

Figure 2.19: Supplementary moments

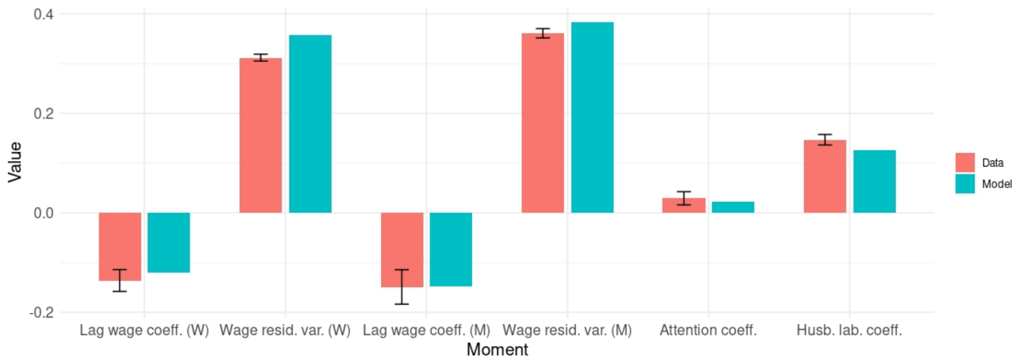

Notes: error bars represent 95% confidence intervals

The fit here is generally good. Notably, the model is able to match closely how agents' labour supply changes when they become attentive, suggesting that the model is capturing the degree to which agents have to change their behaviour once they realise they have made a mistake in their pension planning. The only notable misses are that earnings in the model are slightly more volatile than they are in the data.

In Figure 2.48 in Appendix 2.B.5 I compare model fit for women's labour supply for agents who start the model inattentive in the baseline model versus an alternative model where agents are always attentive. The main finding is that assuming a fully attentive model does not notably reduce the model fit. As such, while accounting for rational inattention might be important for welfare analysis, particularly given the heterogeneous welfare impacts of mistakes, the presence or absence of such mistakes in the model does not make significant changes to aggregate moments.

2.5 Welfare analysis and counterfactuals

2.5.1 How do people change their labour supply when they realise their mistakes?

With the estimated model in hand it is possible to generate truly exogenous shocks to agents' beliefs and analyse how they respond in terms of their labour supply, without the endogeneity issues discussed in Section 2.2.3. This allows us to assess how people change their labour supply when they realise they have made mistakes in their retirement planning.

To this end, I simulate the labour supply of agents over time from age 50 in 2006 to age 70 in 2026, and who are facing a SPA of 65, under three separate assumptions: a) all agents are aware throughout of the true SPA, b) all agents are unaware until the age of 55, when all become aware and c) all agents are unaware until the age of 60, when all become aware. Aware agents are certain of the true SPA but unaware agents have belief g^0 over the SPA taking different values. I then calculate the percentage point difference in labour supply under assumptions b) and c) relative to the labour supply under assumption a). The results are presented in Figure 2.20[26].

Figure 2.20: Probability of working relative to base case

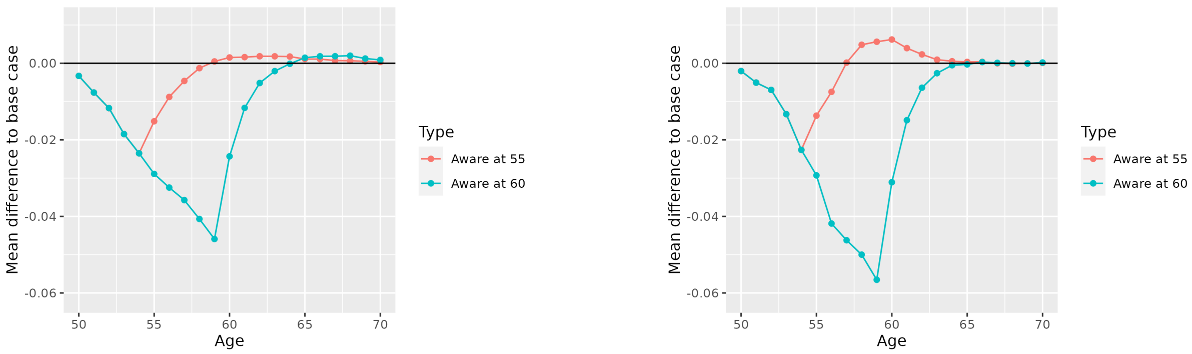

(a) Women

(b) Men

The unaware agents start reducing their labour supply (of the woman in the household and the husband, if any) sooner than the aware agents, because they are over-optimistic about their income in retirement, which creates a widening gap in the labour supply of aware and unaware agents. When agents realise at 55, they increase their labour supply and eventually provide slightly more labour than the agents who were aware from the start. However, the increase in labour supply is not very drastic as they realise 5 years before they could have received any state pension income even under the most optimistic outcome. For those who realise only at age 60, however, the change in labour supply is more dramatic, switching from working 4.3pp less than always-aware agents at age 59 to working around 0.1pp more than the always-aware agents between the ages of 65 and 70. Notably, agents once they become aware do not end up compensating by working much more than always-aware agents in their 60s, which points to the importance of other frictions and wage reductions limiting the appeal of working in one's 60s as a channel of adjustment to the shock. Overall, the model suggests there would be changes in labour supply behaviour after agents realise they have been overoptimistic about their income in retirement, but the changes are substantive only in the cases where agents realise their over-optimism close to retirement.

[26] The corresponding raw figures, without differencing, are presented in Figure 2.49 in Appendix 2.B.6.

2.5.2 Who loses out the most from mistakes in retirement planning?

To analyse the welfare effects of mistakes in retirement planning, I construct a counterfactual sample of agents' labour supply and consumption behaviour, imposing that everyone is always aware of the reform. In other words, in 2006 when the simulation starts, all agents are aware of the reform. I then calculate the difference in experienced labour and consumption utility for the counterfactual agents versus the agents in the baseline model[27].

To provide measures of the utility losses from mistaken beliefs, I take two approaches. For the first approach, to find a money-metric measure of utility losses from mistaken beliefs, I calculate for each inattentive agent in the data how much more wealth they would have to start the simulation with to be indifferent ex post between this and starting off the model attentive with the original amount of wealth, taking into account only experienced utility from consumption and labour. I then interpret this required compensating wealth as the money-metric welfare loss associated with having mistaken beliefs. For the second approach, to find a consumption-metric measure of utility losses from mistaken beliefs, I find the across-the-board increase in consumption per period required for households who start the model inattentive so that they are indifferent ex post between starting the model inattentive with higher consumption per period and starting the model attentive.

The average money-metric welfare loss across all inattentive households was £457, implying an average welfare loss across all households (inattentive or attentive) of £95. The consumption-metric measure of welfare loss is such that starting the model attentive is worth a 0.25% increase in per-period consumption per inattentive household. To put these figures into context, median household consumption at age 60 in the model is £20.3k every two-year period, so the welfare cost of inattention is small relative to the resources of the typical household. This welfare loss is also small relative to the Parliamentary and Health Service Ombudsman's recommendation of compensation of £1000-£2950 to affected women (PHSO 2024)[28].

However, this headline figure masks some notable heterogeneity in how much people would benefit from the elimination of inattention, depending on their demographics. To

[27] I focus on experienced utility from labour and consumption because utility from attention - i.e. attention preference shocks, and the costs incurred from doing one's own research - is mechanically linked to the presence of inattention, and insofar as attention preference shocks can be interpreted both as preference shocks but also randomness in information received, it is not straightforward to interpret the ensuing increase in utility from removing inattention. What I am interested in instead is the deviation from optimal consumption and labour choices which comes from the existence of inattention.

[28] PHSO (2024) does not specify how women should be identified as being affected, recognising that the DWP's maladministration will have had heterogeneous effects, that some women in the relevant cohorts will not have suffered an injustice, that compensating all women in the relevant cohorts would be very expensive but that "finite resources should not be used as an excuse for failing to provide a fair remedy" (PHSO 2024, p. 91).

examine this, I run an OLS regression of money-metric welfare losses on state variables in the first period for agents who start the model inattentive. The results are shown in Figure 2.21.

Figure 2.21: Money-metric welfare loss from having mistaken beliefs

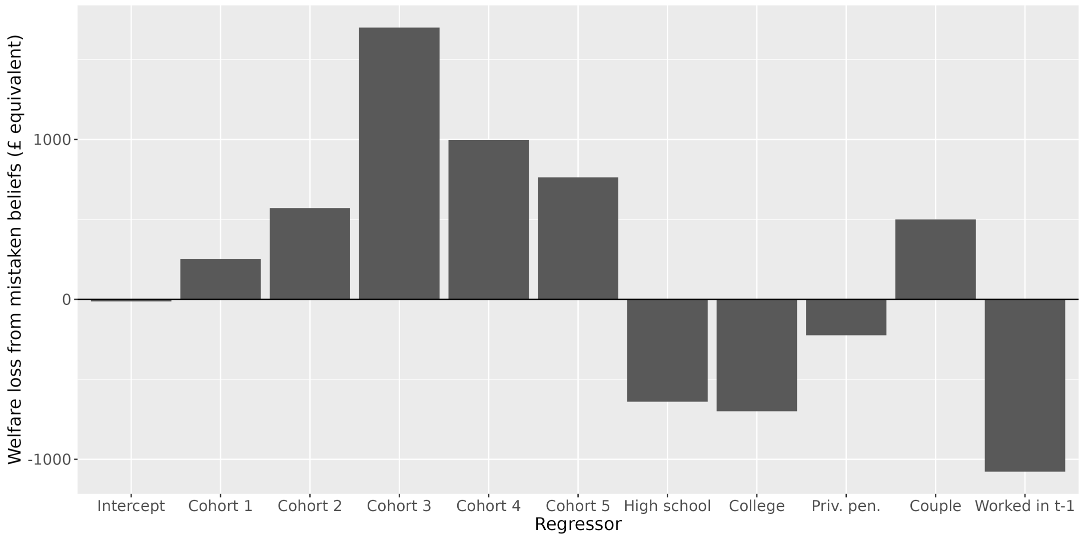

Notes: inattentive households in 2006 only.

The chart shows that the cohort that lost out the most from the imperfect communication of the reform is cohort 3, i.e. those who were born in 1951 or 1952 and thus turned 54 or 55 in 2006. Even though this cohort did not see the biggest increase in the SPA, they suffer the most from inattention because people in this cohort will become attentive closer to the SPA so will have to make more radical changes to their consumption and labour supply behaviour.

Unsurprisingly, more educated people benefit less from the elimination of inattention, because their higher wage means that it is easier for them to cover any unexpected shortfall in state pension income. Likewise, those with private pensions benefit less from the elimination of inattention due to the increased protection of a private pension.

It is somewhat surprising that couples benefit more from the elimination of inattention than singles do, insofar as couples have two income streams so are better protected against unexpected shocks to pension income. However, because there are two agents in a couple there is more scope to optimally adjust behaviour when attentive, which means that couples benefit (slightly) more from the elimination of inattention than singles do.

Finally, those people with more attachment to the labour force benefit less from the elimination of inattention. These people are better protected against shocks to state pension income as they can simply extend their stays in the labour force while keeping consumption relatively constant, whereas those who have already left the labour force may find it prohibitively costly to return and thus need to cut consumption to make ends meet.

2.5.3 Policies to mitigate the effects of mistakes in retirement planning

The problem of agents making mistakes in their retirement planning points to a possible role for the government in mitigating the effects of these mistakes. In this section I use the model to examine some potential policy responses.

2.5.3.1 Wage subsidy for re-entering work

One policy response would be to reduce the costs of re-entering the labour force. Inattention is costly only insofar as the consequences of decisions made when inattentive cannot be reversed when the agent becomes attentive, and a key reason for this might be that there are significant wage and utility costs associated with re-entering the labour market. This is reflected in the results of Figure 2.21, which show that the costs of inattention are much higher for those who were not working at the start of the model, hence have less ability to correct their previous mistakes by extending their working life.

To this end, I consider a very simple wage subsidy of £1000 for those starting work again having not worked the previous period. I suppose that this policy was in place in 2006, and then calculate the new cost of inattention - i.e. what inattentive agents would have been willing to pay, ex post, to start the model attentive - in this counterfactual policy environment.

The wage subsidy has only limited effects in mitigating the cost of inattention, which decreases from its original level of £457 to £393, a 14% reduction. Such a policy is largely ineffective in reducing the costs of inattention because the most significant factor preventing people from re-entering the labour force late in life is not the wage penalty associated with doing so (a penalty of 18% according to the results in Table 2.2) but instead the very significant utility penalty associated with doing so. As discussed in Section 2.4.4, the utility penalty associated with finding work again for the sample under consideration is such that if Agent A was identical to Agent B apart from the fact that Agent A had worked the previous period, and Agent A had a 50% chance of working today, then Agent B would have only a 1.8% probability of working today. As such, relatively small financial incentives to rejoin the labour force do not significantly reduce the extent to which dropping out of the labour force in one's 50s and 60s is an almost permanent decision, and hence do not mitigate the costs of inattention to any great degree.

2.5.3.2 Increase in welfare payments

A separate policy would be to increase welfare payments, with the motivation for such a policy being to protect those people who retire "too early" from the consequences of their mistakes, by ensuring they have a higher income as they wait for their pension to come in. I simulate the effects of such a policy by increase non-labour income for people below the age of 60 who are not working. For ease of comparison with the wage subsidy of the previous section, I calibrate the increase in welfare payments so that it has the same cost to the exchequer as the wage subsidy, resulting in an annual increase in welfare payments for the under-60s of £541[29].

The mitigation of the costs of inattention is again limited. In this counterfactual policy environment the cost of inattention is £420, only an 8% reduction. The issue here is that although the increase in welfare payments does successfully increase the consumption of those on lowest incomes, the benefit is far less targeted than the wage subsidy when it comes to benefitting those that are inattentive. In particular, 10% of those women in receipt of the higher welfare payment were inattentive compared to 19% of those women receiving the wage subsidy for re-entering work. As such, the merits of the welfare payment boost in terms of mitigating the cost of inattention are limited.

2.5.3.3 Increase in information provision

A final policy considered here for the mitigation of the effects of mistakes in retirement planning is an increase in information provision on the part of the government. For instance, as part of a campaign to increase awareness of the SPA increases, the government sent out letters to affected women (PHSO 2021). Here I analyse the effects of information treatments in the abstract, in particular an information treatment bringing about an unconditional increase of 1% in the probability of an inattentive agent becoming attentive in any period.

An increase of 1% in the probability of becoming attentive in any given period results in a reduction of the cost of inattention of £3.93, and therefore (averaging over all households, including attentive ones) a benefit of £0.77 per household. As such, we can state a policy like the letter-writing campaign would offer value for money provided that the cost of supplying the treatment is less than £0.77 per household and it had an effect on awareness at least as positive as a 1% increase in the probability of becoming aware per period.

While information on the cost of the government's SPA letter-writing campaign or its effects on awareness is not available, a non-personalised 16-page leaflet on the 2016 UK EU referendum sent by the UK government to 27 million households cost £9.3 million, or £0.34 per household, to produce, disseminate, and promote via a website (BBC 2026). Given this, it is plausible that a SPA letter-writing campaign would cost less than £0.77

[29]This figure comes from observing that in the counterfactual simulations for the wage subsidy 6.7% of all people in the sample were receiving the wage subsidy but 12.4% would be in receipt of the boost to welfare payments, hence $6.7/12.4 = 0.541$ for the calculation of the welfare payment boost. Note that this abstracts from any behavioural responses to the welfare payment boost in terms of there being less incentive to work.

per household, so assuming it would have a stronger effect on awareness than a 1% increase in the probability of becoming aware every period, we can conclude that such a campaign would likely offer value for money.

2.6 Conclusion

In this paper I offer an assessment of the labour supply and welfare consequences of mistakes in retirement planning, exploiting data on agents' mistaken beliefs about the age at which they can receive their state pension. I find that agents with mistaken beliefs suffer an average cost of £457, or 0.25% of per-period consumption, from having those mistaken beliefs, pointing to notable limitations in agents' ability to costlessly change their behaviour to correct their previous mistakes. The costs of mistaken beliefs are particularly high for those agents closest to retirement. Once agents realise their mistakes, they increase their labour supply to compensate for their loss of expected income. However, ultimately, the welfare losses are small relative to the resources of the median household close to retirement.

While the current paper has included several possible channels through which mistakes in retirement planning can be costly, such as utility and wage costs of re-entering the labour force having not worked in the previous period, or increasing disutility/decreasing wages of agents as they grow older, it has abstracted away from other important reasons why such mistakes are costly. Incorporating more sophisticated models of human capital depreciation, health and caring responsibilities, and non-separability in consumption and labour utility would allow a more in-depth assessment of the costs of mistakes in retirement planning, while allowing for non-unitary decision-making within the household would allow for study of who bears the burden of these mistakes. I leave these questions to future research.

References

Banks, James, Richard Blundell, et al. (2012). "Housing Mobility and Downsizing at Older Ages in Britain and the USA". In: *Economica*.

Banks, James, Carl Emmerson, et al. (2005). *Prepared for retirement? The adequacy and distribution of retirement resources in England*. Research Report R67. ISBN: 9781903274439. IFS Report. DOI: 10.1920/re.ifs.2005.0067.

Battistin, Erich et al. (2009). "The Retirement Consumption Puzzle: Evidence from a Regression Discontinuity Approach". In: *American Economic Review* 99.5, pp. 2209–26. ISSN: 0002-8282. DOI: 10.1257/aer.99.5.2209.

BBC (2026). *EU referendum: PM 'makes no apology' for £9m EU leaflets.* URL: https://www.bbc.co.uk/news/uk-politics-eu-referendum-35984991 (visited on 02/19/2026).

Blundell, Richard, Rowena Crawford, et al. (2016). "Comparing Retirement Wealth Trajectories on Both Sides of the Pond". In: *Fiscal Studies.*

Blundell, Richard, Eric French, and Gemma Tetlow (2016). "Chapter 8 - Retirement Incentives and Labor Supply". In: *Handbook of the Economics of Population Aging.* Ed. by John Piggott and Alan Woodland. Vol. 1. North-Holland, pp. 457–566. DOI: 10.1016/bs.hespa.2016.10.001.

Brown, Zach Y. and Jihye Jeon (2024). "Endogenous Information and Simplifying Insurance Choice". In: *Econometrica* 92.3, pp. 881–911. ISSN: 1468-0262. DOI: 10.3982/ECTA18555.

Crawford, Rowena and Cormac O'Dea (2020). "Household portfolios and financial preparedness for retirement". In: *Quantitative Economics* 11.2, pp. 637–670. ISSN: 1759-7331. DOI: 10.3982/QE725.

Cribb, Jonathan, Carl Emmerson, and Gemma Tetlow (2016). "Signals matter? Large retirement responses to limited financial incentives". In: *Labour Economics* 42, pp. 203–212. ISSN: 0927-5371. DOI: 10.1016/j.labeco.2016.09.005.

De Nardi, Mariacristina, Eric French, and John B. Jones (2010). "Why Do the Elderly Save? The Role of Medical Expenses". In: *Journal of Political Economy* 118.1. Publisher: The University of Chicago Press, pp. 39–75. ISSN: 0022-3808. DOI: 10.1086/651674.

De Nardi, Mariacristina, Eric French, and John Bailey Jones (2016). *Savings after Retirement: A Survey.* Rochester, NY. DOI: 10.1146/annurev-economics-080315-015127.

Department for Work and Pensions (2013). *The single tier pension: a simple foundation for saving.* URL: https://assets.publishing.service.gov.uk/media/5a7b622940f0b64646935cb1/single-tier-pension.pdf (visited on 05/12/2025).

— (2014). *State Pension age timetable.* Tech. rep.

— (2024). *Government response to the Parliamentary and Health Service Ombudsman's investigation into Women's State Pension age and associated issues.* URL: https://www.gov.uk/government/publications/government-response-to-parliamentary-and-health-service-ombudsmans-investigation-into-womens-state-pension-age-communications-and-associated-issues (visited on 05/12/2025).

— (2025a). *Benefit and pension rates 2025 to 2026.* URL: https://www.gov.uk/government/publications/benefit-and-pension-rates-2025-to-2026/benefit-and-pension-rates-2025-to-2026#state-pension (visited on 05/12/2025).

— (2025b). *Your state pension explained.* URL: https://www.gov.uk/government/publications/your-new-state-pension-explained/your-state-pension-explained (visited on 05/12/2025).

Department of Health (2011). *Fairer Care Funding.* Tech. rep. Last Modified: 09/09/2011 17:24 Publisher: Department of Health, Richmond House, 79 Whitehall, London SW1A 2NJ, UK, dhmail@dh.gsi.gov.uk.

Deshpande, Manasi, Itzik Fadlon, and Colin Gray (2024). "How Sticky Is Retirement Behavior in the United States?" In: *The Review of Economics and Statistics* 106.2, pp. 370–383. ISSN: 0034-6535. DOI: 10.1162/rest_a_01151.

García-Miralles, Esteban and Jonathan M. Leganza (2024a). "Public Pensions and Private Savings". In: *American Economic Journal: Economic Policy* 16.2, pp. 366–405. ISSN: 1945-7731. DOI: 10.1257/pol.20220019.

— (2024b). *Savings responses to increasing retirement ages.* CEPR. URL: https://cepr.org/voxeu/columns/savings-responses-increasing-retirement-ages (visited on 04/15/2025).

Hentall-MacCuish, Jamie (2025). *Costly attention and retirement.* IFS Working Paper. The IFS. DOI: 10.1920/wp.ifs.2024.5924.

Holman, Daniel, Liam Foster, and Moritz Hess (2020). "Inequalities in women's awareness of changes to the State Pension Age in England and the role of cognitive ability". In: *Ageing & Society* 40.1, pp. 144–161. ISSN: 0144-686X, 1469-1779. DOI: 10.1017/S0144686X1800082X.

House of Commons Library (2024). *The old State Pension.* URL: https://researchbriefings.files.parliament.uk/documents/CBP-10114/CBP-10114.pdf (visited on 05/12/2025).

Kopecky, Karen and Richard Suen (2010). "Finite State Markov-Chain Approximations to Highly Persistent Processes". In: *Review of Economic Dynamics.* DOI: 10.1016/j.red.2010.02.002.

Lalive, Rafael and Pierpaolo Parrotta (2017). "How does pension eligibility affect labor supply in couples?" In: *Labour Economics* 46, pp. 177–188. ISSN: 0927-5371. DOI: 10.1016/j.labeco.2016.10.002.

Lockwood, Lee M. (2018). "Incidental Bequests and the Choice to Self-Insure Late-Life Risks". In: *American Economic Review* 108.9, pp. 2513–2550. ISSN: 0002-8282. DOI: 10.1257/aer.20141651.

Mackley, Andrew, Djuna Thurley, and Roderick McInnes (2021). *Increases in the State Pension age for women born in the 1950s.* Tech. rep.

Nakazawa, Nobuhiko (2022). "The Effects of Increasing the Eligibility Age for Public Pension on Individual Labor Supply: Evidence from Japan". In: *Journal of Human Resources.* Publisher: University of Wisconsin Press Section: Articles. ISSN: 0022-166X, 1548-8004. DOI: 10.3368/jhr.0421-11627R1.

O'Dea, Cormac and David Sturrock (2018). "Subjective expectations of survival and economic behaviour". In: *IFS Working Papers.* Number: W18/14 Publisher: Institute for Fiscal Studies.

— (2020). *Survival Pessimism and the Demand for Annuities.* Rochester, NY.

OECD (2019). *Pensions at a Glance 2019*. OECD. Tech. rep.

Olafsson, Arna and Michaela Pagel (2018). *The Retirement-Consumption Puzzle: New Evidence from Personal Finances*. DOI: 10.3386/w24405.

PHSO (2021). *Women's State Pension age: our findings on the Department for Work and Pensions' communication of changes*. GOV.UK. (Visited on 04/15/2025).

— (2024). "Women's State Pension age: our findings on injustice and associated issues". In.

Royal London (2023). *State Second Pension Explained*. URL: https://adviser.royallondon.com/technical-central/pensions/state-benefits-pension-manuals/state-second-pension-explained/ (visited on 05/12/2025).

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun (2006). "Are Americans Saving "Optimally" for Retirement?" In: *Journal of Political Economy* 114.4. Publisher: The University of Chicago Press, pp. 607–643. ISSN: 0022-3808. DOI: 10.1086/506335.

Staubli, Stefan and Josef Zweimüller (2013). "Does raising the early retirement age increase employment of older workers?" In: *Journal of Public Economics* 108, pp. 17–32. ISSN: 0047-2727. DOI: 10.1016/j.jpubeco.2013.09.003.

Appendices

2.A The UK state pension

2.A.1 Policy context

The state pension is a regular payment from the government which can be claimed by those above state pension age provided they have made or are credited with sufficient National Insurance (NI) contributions.

For those reaching SPA before 6 April 2016[30], the relevant pension system was the "old" state pension (House of Commons Library 2024). This state pension had two tiers - the basic state pension (BSP) and the additional state pension. The basic state pension was a flat-rate benefit which people could build entitlement to by accruing NI qualifying years, i.e. years where they had either paid or been credited with NI contributions. On reaching the full number of qualifying years for the state pension, people qualified for the full state pension, otherwise they qualified for an amount proportional to their qualifying years.

The additional state pension had an earnings-related component. The exact nature of this component has varied over the decades, with the earnings-related component being introduced to the state pension system in 1961 and undergoing important reforms in 1978 and 2002. In particular, between 2002 and 2016, the relevant earnings-related component was the "second state pension", which gave extra payments to those on low and middle incomes (Royal London 2023).

For those reaching SPA after 6 April 2016, the relevant pension system was the "new" state pension (Department for Work and Pensions 2025b). A major objective of the reform was to increase clarity and transparency regarding pension entitlement by transforming it into a single-tier pension (Department for Work and Pensions 2013). Entitlement under the "new" state pension is generally based on NI contributions only, rather than having an earnings-related component (Department for Work and Pensions 2025b). For those people who built up entitlement under the pre-2016 system but had not reached their SPA in 2016, there were transitional arrangements in place under which their previous contributions would supplement the amount they were entitled to under the post-2016 system (Department for Work and Pensions 2013).

2.A.2 Imputing SPAs

In the public version of ELSA, respondents' months of birth are not made available. Instead, there is data on age (in years, at time of interview), year of birth, and month

[30] This would be women born before 6 April 1953 and men born before 6 April 1951.

and year of interview.

To impute month of birth, I calculate for each individual in the data the range of birth months it is possible for them to have, consistent with their age at the time of interview, the month and year of the interview and their year of birth. For instance, if someone is born in 1950 and is 59 years old when interviewed in October 2010, their birth month must be in October (later), November or December of 1950. If respondents are interviewed in different months across waves, more restrictions can be placed on the set of possible birth months. Then, I randomly assign each individual a birth month out of all the birth months which it would be possible for them to have.

Due to the fact that there will be errors in the imputed birth months, I choose a definition of correctness of beliefs in the main text (i.e. less than 2 years out from the truth) that will avoid giving false positives for mistakes.

2.B Supplementary tables and figures

2.B.1 Descriptive statistics

Descriptive statistics on the estimation sample vs. the full ELSA sample from Waves 2 to 9 are presented in Table 2.3, for the women in the two samples.

Table 2.3: Descriptive statistics

	Full ELSA sample	Estimation sample
Age	66.46	60.74
Married	0.61	0.71
Graduated HS	0.48	0.57
College	0.15	0.17
In FT work	0.31	0.37
Has priv. pen.	0.69	0.72
Homeowner	0.86	0.85
N	43406	9235

Notes: table displays means for each variable. Individuals considered to have a private pension if they are ever observed contributing to or receiving a private pension. Means weighted by ELSA person-level weights.

2.B.2 Labour supply by cohort analysis

In Figure 2.3, it is shown that for women labour supply in their early 60s was higher for cohorts with a later SPA (i.e. those born after 1950). Part of the explanation for this may be mechanical in the sense that employment contracts or normal retirement ages might be linked to the SPA.

It is difficult to analyse this issue directly using ELSA data because information on employment contracts is quite scarce, and the key premise of this paper is that people might be incorrect in their beliefs about pension ages, so the information that we do have in ELSA may not be very accurate. However, there is data on people's private pensions and in particular whether each pension has an associated normal retirement age, so it is at least possible to check the extent to which this aspect of employment practices changes with the SPA, at least assuming that agents are aware of their pension scheme's rules.

Figure 2.22 shows the mean normal retirement age by cohort, for women with a private pension scheme with an attached normal retirement age. The first point to note is that broadly speaking younger cohorts born later face later normal retirement ages in their pension schemes.

Figure 2.23 shows a histogram of mean normal retirement ages for women whose state pension age is between 61 (and 0 months) and 64 (and 0 months). Here, the key point is that normal retirement ages are unlikely to take on values that are not 60 or 65 even

Figure 2.22: Mean normal retirement age by cohort

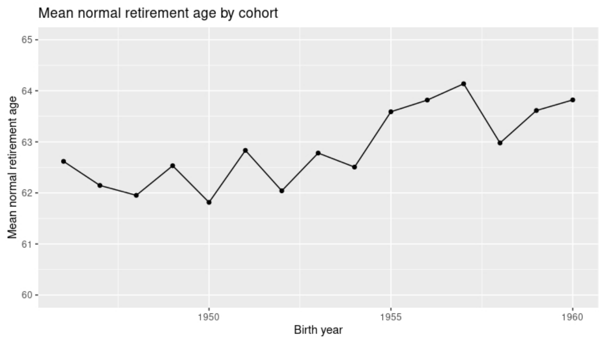

Notes: pensions with normal retirement ages less than 50 are dropped. Means weighted by ELSA individual-level weights.

Figure 2.23: Normal retirement age histogram

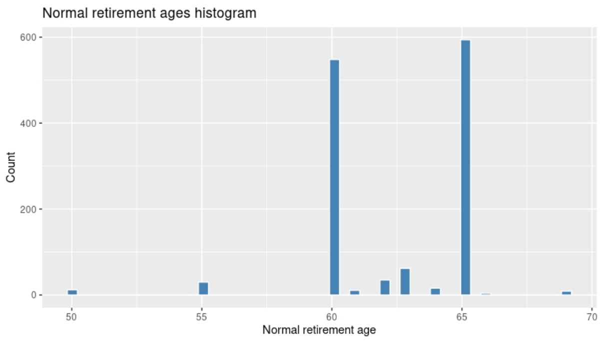

Notes: pensions with normal retirement ages less than 50 are dropped. Histogram presented for those women whose SPA is between 61 and 64.

for those with SPAs between those two values. In other words, even though normal retirement ages tended to increase this does not seem to be because of them being set exactly equal to the SPA by default.

As such, while part of the cohort effects in Figure 2.3 is likely due to increasing normal retirement ages of private pension schemes and other aspects of employment contracts, there is little evidence of one particular aspect of employment contracts around retirement - namely, normal retirement ages - changing mechanically with the SPA.

2.B.3 Correlates of having correct beliefs

Table 2.4: Correlates of having correct beliefs

Dependent Variable:	Correct belief
College	-0.0358
	(0.0298)
Couple	-0.0436
	(0.0273)
In labour force	0.0328
	(0.0245)
Homeowner	0.1136***
	(0.0337)
Long-term planner	0.0529*
	(0.0249)
Socially isolated	-0.0146
	(0.0188)
Financial resp.	-0.0143
	(0.0250)
Has private pension	0.0899***
	(0.0277)
Fixed-effects	
Age	Yes
YOB	Yes
Fit statistics	
Observations	1,470
R^2	0.18914
Within R^2	0.03284

Notes: correct beliefs defined as beliefs within two years of the true SPA. Regression weights are ELSA person-level weights.

2.B.4 Extra event study graphs

2.B.4.1 Correctness of beliefs by time and by age

Figure 2.24 shows the proportion of women who have correct beliefs, by age and by year, across different 2-year birth cohorts.

Figure 2.24: Proportion of women with correct beliefs, by cohort

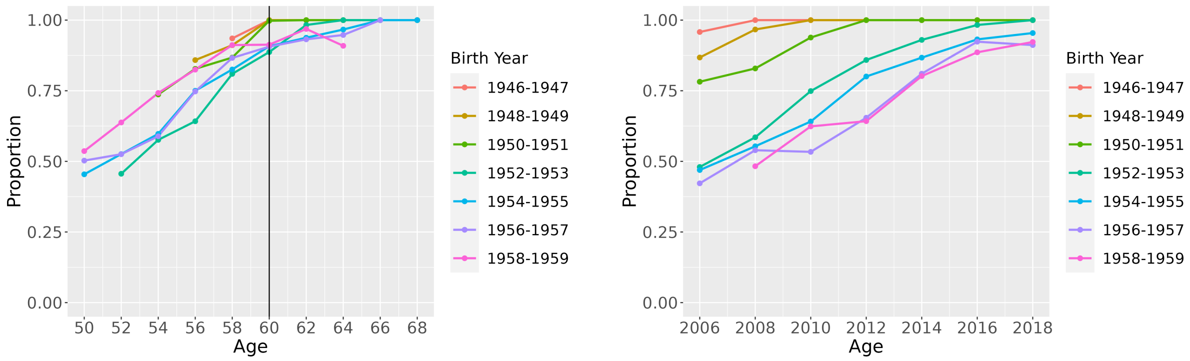

(a) By age (b) By year

Notes: means weighted by ELSA person-level weights.

2.B.4.2 Raw labour force participation graphs

The Figure below presents the equivalent of Figure 2.9 but comparing agents who had correct beliefs in 2006 versus those who did not, across calendar time.

Figure 2.25: Labour force participation by awareness in Wave 3 (2006)

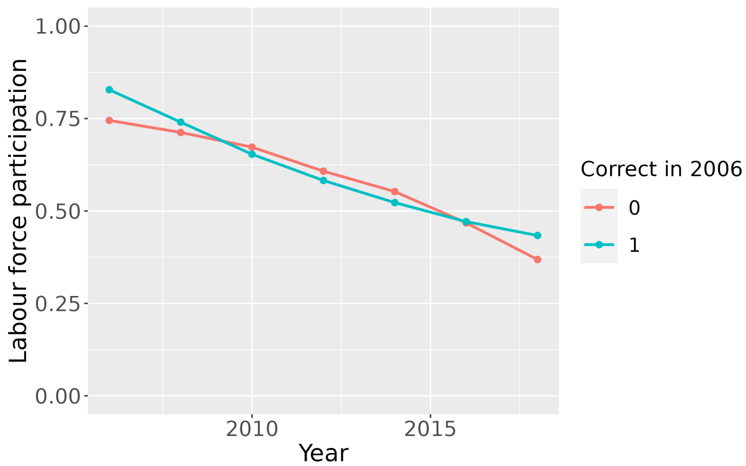

Notes: individuals counted as being in the labour force if self-describe as working full-time, working part-time, or unemployed. Coefficients plotted are net of year-of-birth controls, using the intercepts for the 1952 birth cohort. Means weighted by ELSA person-level weights.

2.B.4.3 Retirement age regression

In Table 2.5 I show the results of a regression with the outcome variable being dummies for being observed retired at various ages on the key RHS variable of whether the individual in question was observed with correct beliefs in 2006, the first wave where the question about SPA beliefs was asked, along with other controls. The results suggest that even controlling for observables those people with correct beliefs in 2006 were more likely to be retired at age 60 and at age 65 than those who did not have correct beliefs.

Table 2.5: Retirement age by correctness of beliefs in 2006

	Ret. at 60	Ret. at 60	Ret. at 65	Ret. at 65
Correct in 2006	0.113***	0.124***	0.128***	0.112**
	(0.033)	(0.032)	(0.036)	(0.037)
YOB dummies	Y	Y	Y	Y
Demographic controls	N	Y	N	Y
Num.Obs.	1447	1351	1083	1009
R2 Adj.	0.034	0.115	0.150	0.170
Mean dep. var.	0.403	0.401	0.654	0.653

Notes: women born between 1945 and 1956, who were present in the sample in Wave 3 (2006) and who were observed at the ages of 60/61 (for Columns 1 and 2) or 65/66 (for Columns 3 and 4). "Ret. at 60" is a dummy equal to 1 if the respondent is observed at age 60/61 and self-identifies as retired at that age, and 0 otherwise (due to the biennial nature of ELSA, two-year age bins must be used). "Ret. at 65" is a dummy equal to 1 if the respondent is observed at age 65/66 and self-identifies as retired at that age, and 0 otherwise. "Demographic controls" are dummies for education status, couple status, whether in the labour force, and occupation if working, with these demographic controls calculated for 2006. Regression weights are ELSA person-level weights.

2.B.4.4 Alternative outcome variables

Figure 2.26 shows the results of the event study regression in Equation 2.1 where the outcome variable is a dummy for rejoining the labour force.

Figure 2.26: Event study: proportion rejoining labour force

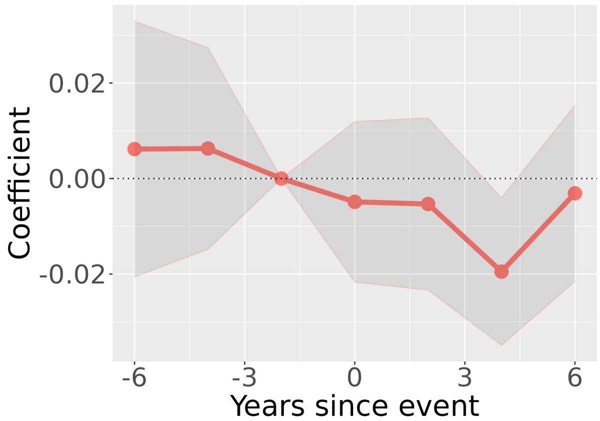

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.27 shows the results of the event study regression in Equation 2.1 where the outcome variable is a dummy for working full-time or part-time (i.e. rather than being in the labour force more generally).

Figure 2.27: Event study: proportion working

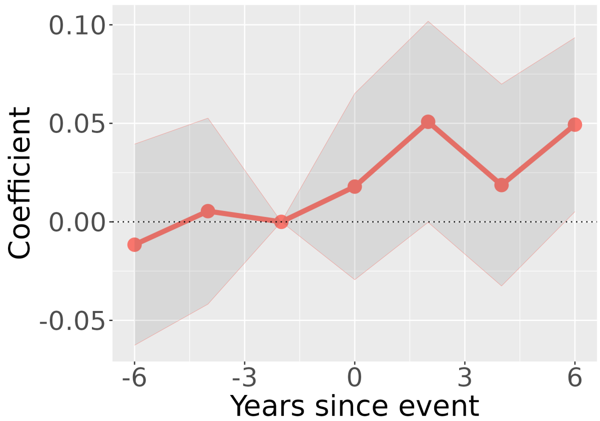

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

2.B.4.5 Event studies with stricter definition of having correct beliefs

The figures below are the equivalents of the event studies in Section 2.2.3 except that in these event studies a person is counted as switching to have correct beliefs at time t iff she is observed with incorrect beliefs at $t-1$ and correct beliefs at t and $t+1$, thus ruling out as moments of realisation cases where agents had correct beliefs for one period but then went back to having incorrect beliefs.

Figure 2.28: Event study (alternative definition of correct beliefs): labour force participation

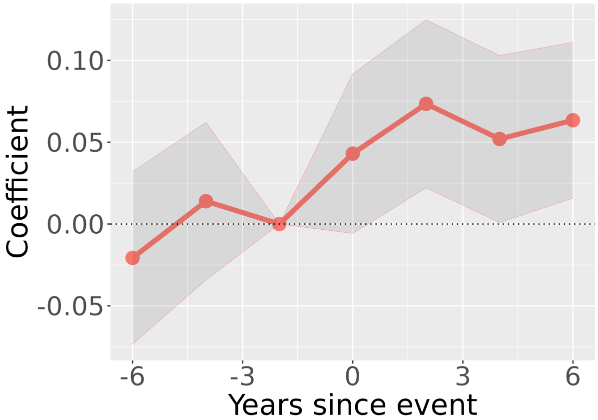

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.29: Event study (alternative definition of correct beliefs): proportion leaving labour force

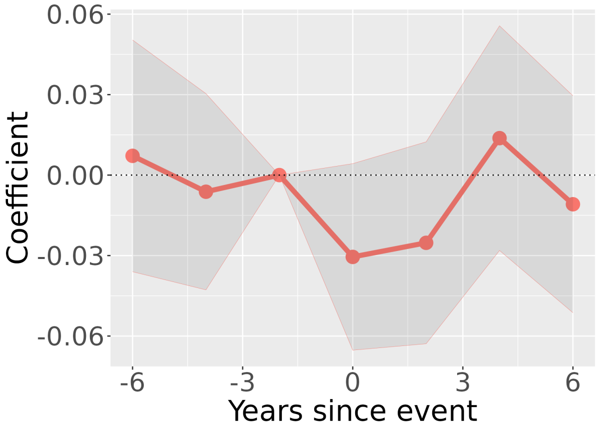

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.30: Event study (alternative definition of correct beliefs): annual labour earnings

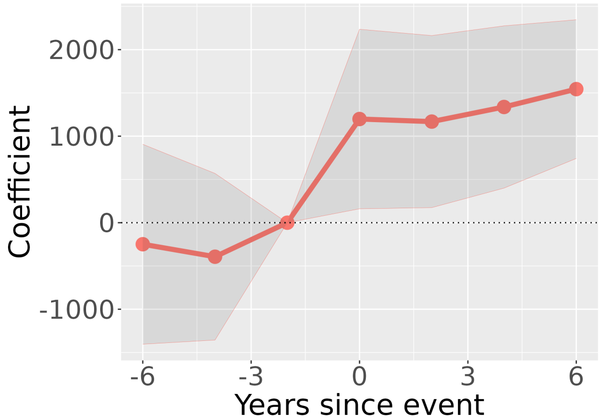

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.31: Event study (alternative definition of correct beliefs): spouse's labour force participation

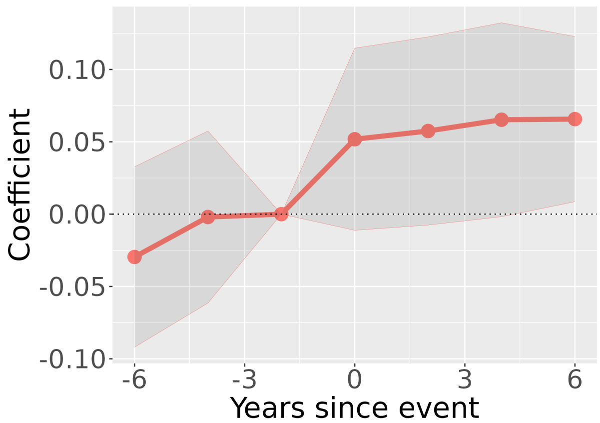

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.32: Event studies (alternative definition of correct beliefs): other outcome variables

(a) Log annual food spending

(b) Own home

(c) >£10k financial wealth

(d) Financially struggling

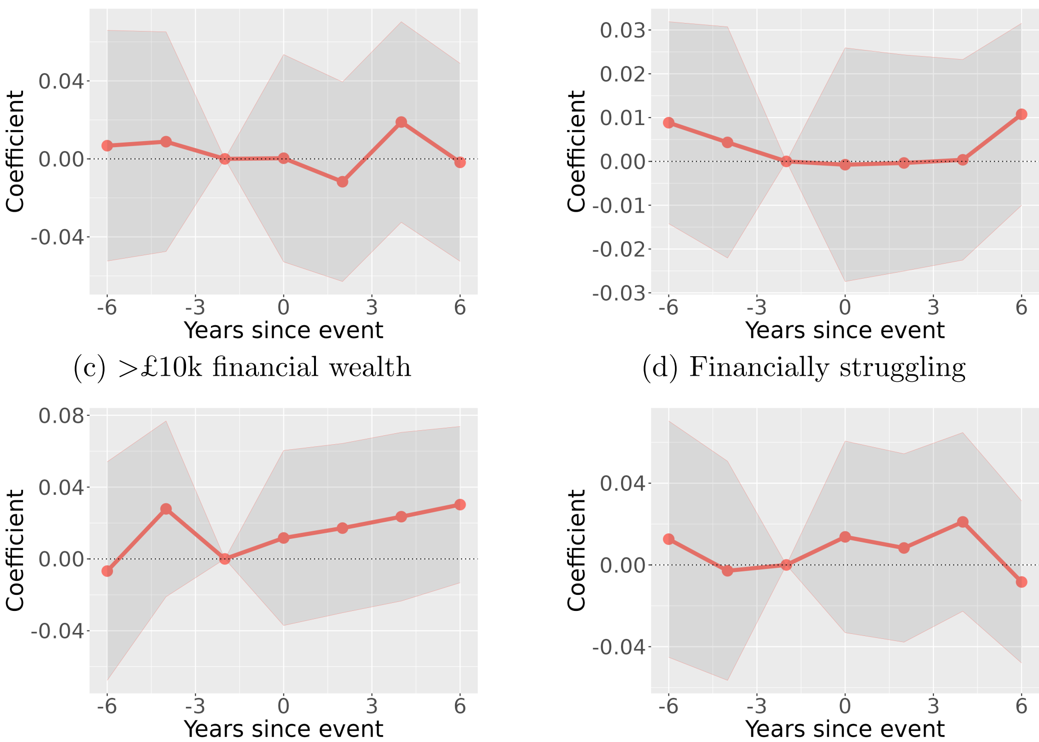

Notes: standard errors clustered at the individual level. Shaded area shows 95% confidence interval.

2.B.4.6 Event studies with added age trends by occupation

The figures below are the equivalents of the event studies in Section 2.2.3 except that in these event studies there are additional trends by occupation status in 2006. Specifically, I take the NS-SEC measure of an individual's occupation included in ELSA, adding an extra category for not working. I then interact these occupation dummies with age and include them in the regressions, to allow for potential differences in age trends across occupations.

Figure 2.33: Event study (added occupation-based age trend): labour force participation

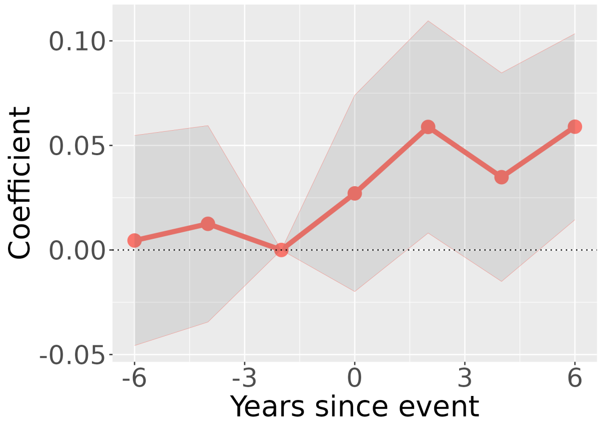

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.34: Event study (added occupation-based age trend): proportion leaving labour force

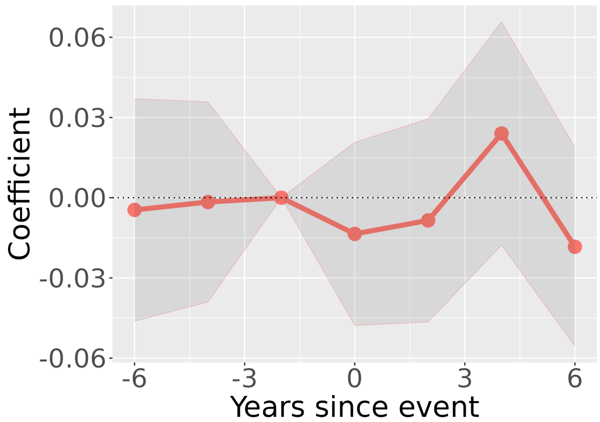

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.35: Event study (added occupation-based age trend): annual labour earnings

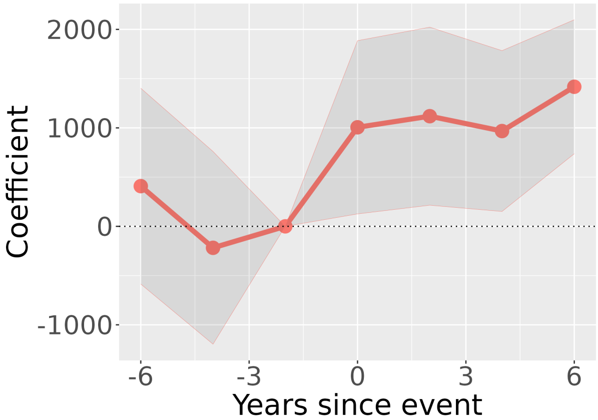

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.36: Event study (added occupation-based age trend): spouse's labour force participation

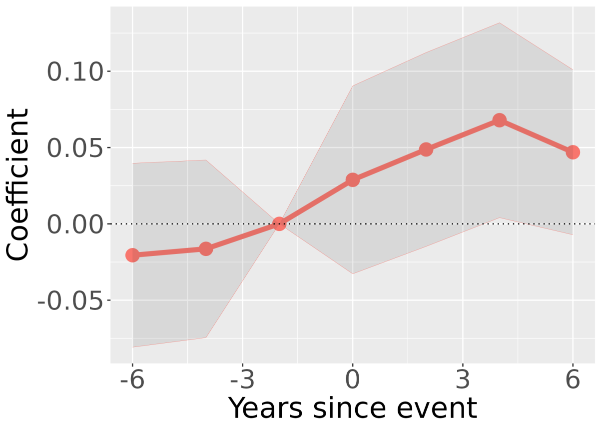

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.37: Event studies (added occupation-based age trend): other outcome variables

Notes: standard errors clustered at the individual level. Shaded area shows 95% confidence interval.

2.B.4.7 Event studies without differential age trends by initial labour force status

The figures below are the equivalents of the event studies in Section 2.2.3 except that in these event studies I remove age trends by labour force status of the woman (and husband, if any) in 2006.

Figure 2.38: Event study (no differential age trend): labour force participation

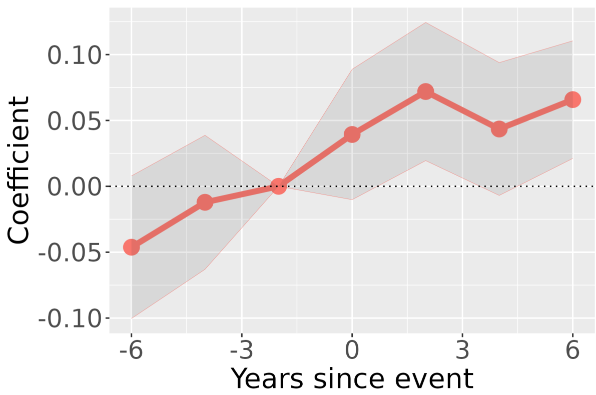

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.39: Event study (no differential age trend): proportion leaving labour force

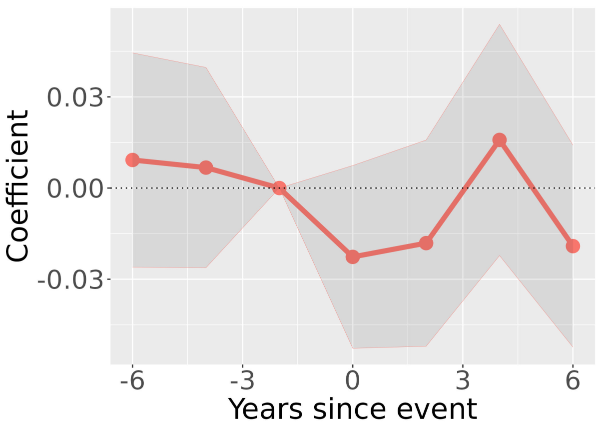

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.40: Event study (no differential age trend): annual labour earnings

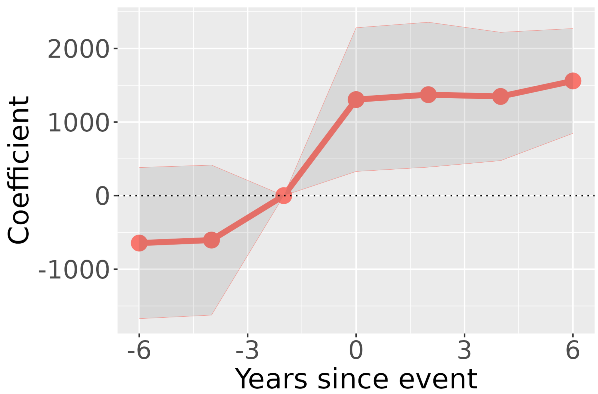

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.41: Event study (no differential age trend): spouse's labour force participation

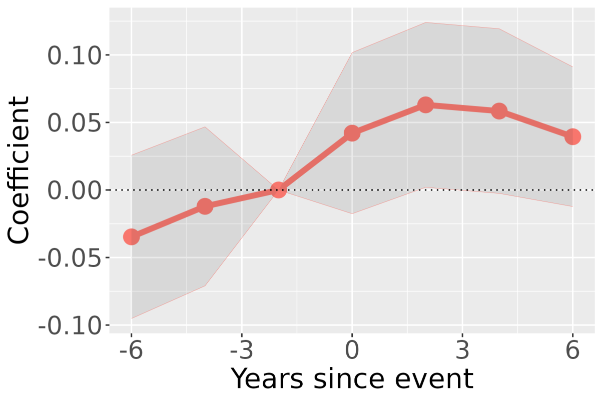

Notes: standard errors clustered at individual level. Shaded area shows 95% confidence interval.

Figure 2.42: Event studies (no differential age trend): other outcome variables

Notes: standard errors clustered at the individual level. Shaded area shows 95% confidence interval.

2.B.5 Extra model fit graphs

2.B.5.1 Model fit breakdown by cohort

These figures break down the model fit graphs of Section 2.4.5 by separating out means by cohort.

Figure 2.43: Model fit - women's probability of working

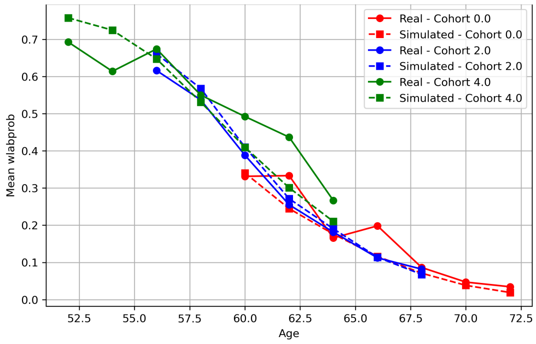

(a) Even-numbered cohorts

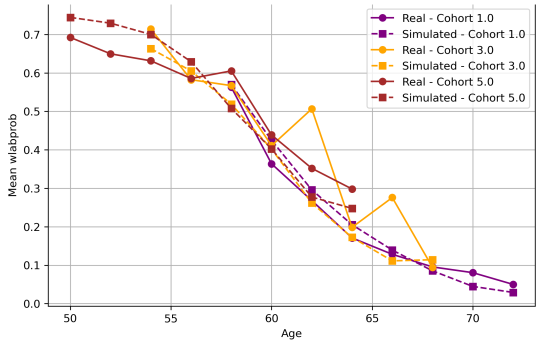

(b) Odd-numbered cohorts

Figure 2.44: Model fit - women's labour force earnings

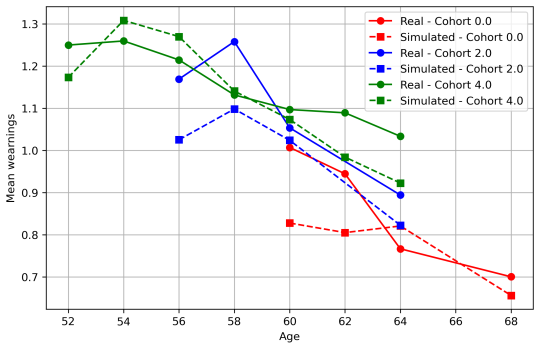

(a) Even-numbered cohorts

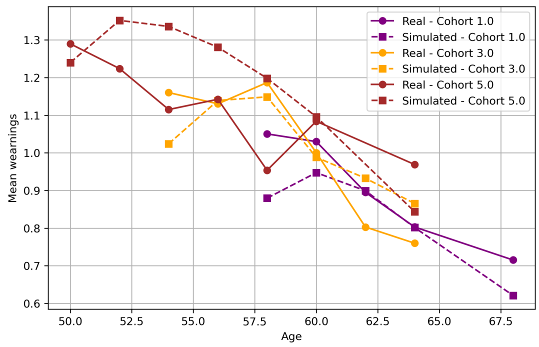

(b) Odd-numbered cohorts

Figure 2.45: Model fit - men's probability of working

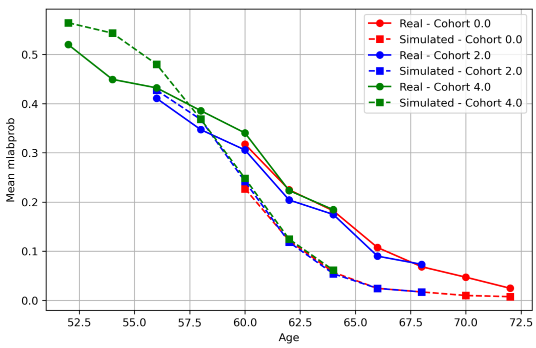

(a) Even-numbered cohorts

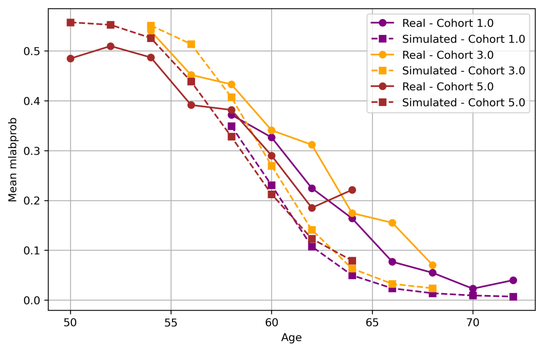

(b) Odd-numbered cohorts

Figure 2.46: Model fit - men's earnings

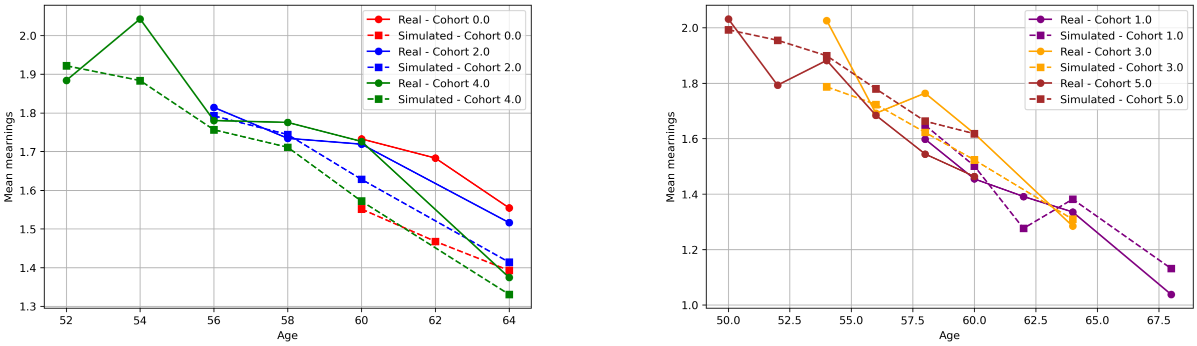

(a) Even-numbered cohorts (b) Odd-numbered cohorts

Figure 2.47: Model fit - proportion attentive

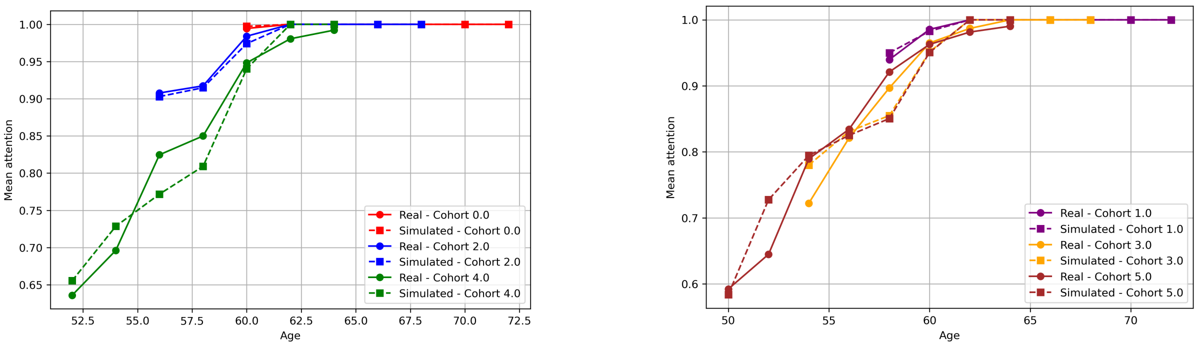

(a) Even-numbered cohorts (b) Odd-numbered cohorts

2.B.5.2 Comparing baseline model with full attention model

The below Figure 2.48 shows the model fit for women's labour supply for women who are inattentive in 2006 in the data, comparing their mean labour supply in the data, their labour supply in the baseline model and their labour supply in an alternative model where all agents are always attentive.

Figure 2.48: Comparison of model fit in baseline model vs. full attention model

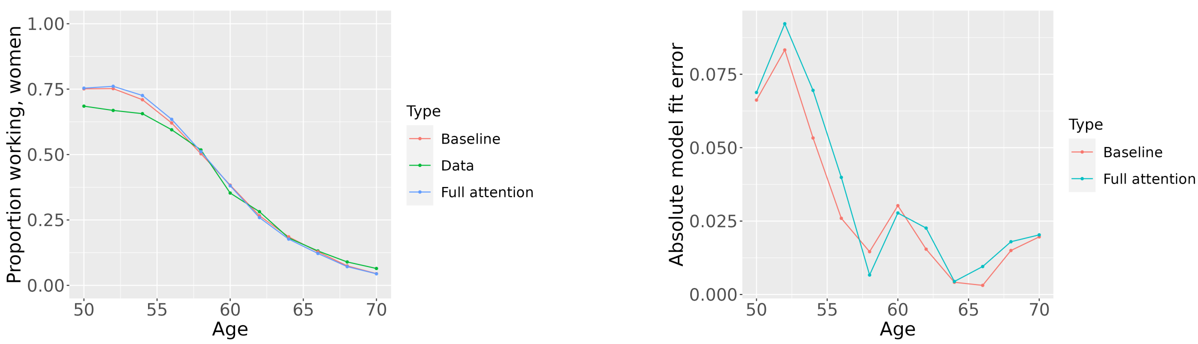

(a) Raw means (b) Absolute difference relative to data

Although the model fit of the baseline model is better, the difference is slight, with the noise in the real data being more important than any gain in model fit between the two models.

2.B.6 Extra simulation graphs

The Figure 2.49 below shows the equivalents of Figure 2.20 but in terms of raw means.

Figure 2.49: Probability of working under different awareness scenarios

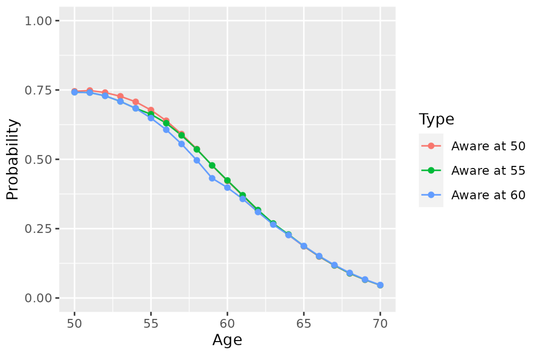

(a) Women

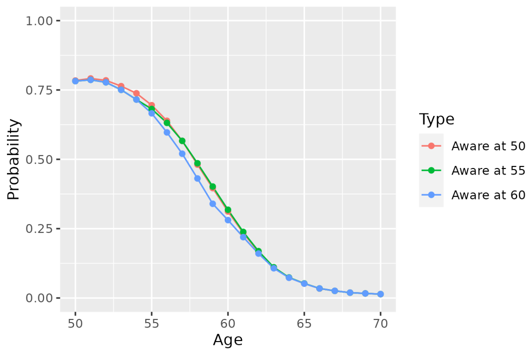

(b) Men

2.B.7 Re-estimation with total wealth

Table 2.6 shows the results of the re-estimation of the model using total wealth, housing plus financial, rather than just financial wealth.

As stated above, the key change is that there is now a smaller penalty for working, which mechanically is because higher wealth drives down labour supply so the cost of working must also decrease in order to keep labour supply at realistic levels.

As for welfare analysis, using the money-metric and consumption-metric approaches outlined in Section 2.5, the cost of inattention in the model with housing wealth is only £83, compared to £457 in the model with financial wealth, which correspond in value to a 0.004% versus 0.25% increase in per-period consumption.

As such, if the relevant wealth variable is total wealth, the costs of inattention are much smaller, though as discussed in the main text it is rare for UK households to access their housing wealth in old age, suggesting that using total wealth as the relevant wealth variable will exaggerate the extent to which households are protected against shocks.

Table 2.6: Estimation Results - total wealth counterfactual

Parameter	Estimate
Work penalty parameters	
ω_0 - Constant	2.266
	(0.041)
ω_{age} - Linear age trend	-0.015
	(0.001)
ω_{age2} - Quadratic age trend	0.002
	(0.001)
$\omega_{prevwork}$ - Worked last period	-2.294
	(0.051)
ω_{joint} - Joint leisure	-0.197
	(0.003)
(Log) Wage parameters	
γ_0 - Constant	-0.398
	(0.025)
γ_{HS} - High school education	0.274
	(0.007)
γ_{col} - College education	0.787
	(0.041)
γ_{age} - Linear age trend	-0.025
	(0.001)
γ_{age2} - Quadratic age trend	-0.001
	(0.000)
$\gamma_{prevwork}$ - Worked last period	0.189
	(0.004)
γ_{male} - Male	0.234
	(0.048)
Attention penalty parameters	
κ_0 - Constant	1.750
	(0.033)
κ_{HS} - High school education	-0.103
	(0.001)
κ_{col} - College education	-0.311
	(0.010)
$\kappa_{prevwork}$ - Worked last period	-0.018
	(0.003)
κ_{couple} - Couple	-0.319
	(0.031)
Other parameters	
$\sigma^2_{\nu w}$ - Variance of iid component of prod. shock, women	0.488
	(0.046)
ρ_w - Persistence of prod. shock, women	0.399
	(0.038)
$\sigma^2_{\nu m}$ - Variance of iid component of prod. shock, men	0.738
	(0.066)
ρ_m - Persistence of prod. shock, men	0.544
	(0.010)
λ - Belief decay parameter	0.349
	(0.048)
σ_ϵ - Type 1 EV spread parameter, labour	0.551
	(0.011)
σ_ξ - Type 1 EV spread parameter, attention	0.473
	(0.054)

Notes: estimation via MSM. See Section 2.4.2 for discussion of moments and weighting matrix used. Standard errors calculated from 20 bootstrap replications. For the calculation of age trends, note that ages were normalised such that age 50 in the data was age 0 in the model.

3 | Long-term care policy and housing market efficiency

Long-term care policy and housing market efficiency

Abstract

In many countries government policy on funding long-term care for older people incentivises holding housing wealth over financial wealth through exempting housing wealth from the test for means-tested government support with long-term care costs (a “homestead exemption”). I analyse the degree to which such exemptions distort the housing demand of older people and the effects on younger people through the housing market using the UK as my setting. I build and estimate an overlapping generations model of the housing market where multiple generations trade houses over the course of their life cycles while facing income, longevity and health risk. By comparing housing market steady states with and without the homestead exemption, I find that a budget-balanced removal of the homestead exemption would reduce house prices by 6.5% and increase welfare by an equivalent of a £144 (0.6%) annual increase in consumption per household. The main beneficiaries are those with less housing wealth in the initial steady state, whereas those who lose out most are those with long-term care problems and more inherited wealth in the initial steady state.

3.1 Introduction

Ageing populations around the world are increasing the demand for government-funded long-term care (LTC) and placing pressure on public finances (Gruber et al. 2025). The sustainability of current state provision of LTC has been called into question in the US (MACPAC 2026), the UK (Nuffield Trust 2025), Canada (Woolley 2023), and many European Union countries (Mosca et al. 2016). Policymakers thus face difficult decisions about how best to fund LTC, with many countries proposing or enacting major reforms to their LTC systems (Dilnot 2011; Yamada et al. 2020).

A separate policy issue that is receiving widespread attention is a burgeoning housing affordability crisis. In many countries the supply of housing cannot keep up with demand for housing so prices rise (Hilber 2023; Lucy 2025; Reisenbichler 2025). In the UK case, rents take up an increasing proportion of household incomes (ONS 2025c) and households take longer to save up to buy a house, and in some cases are priced out of the housing market entirely (National Housing Federation 2025).

A thread which ties these two apparently distinct policy questions together is that of the "homestead exemption" for means-tested support with LTC costs, whereby a household's principal residence is in many cases not counted as part of eligible assets for the means test. Countries like the UK and US provide means-tested long-term care whereby people's wealth is assessed and those falling below a wealth threshold are eligible for state support with their long-term care costs whereas those with wealth above the threshold have to self-fund. While the exact nature of this exemption will differ between countries, and often within countries, the main principle is the same: housing wealth is better protected from being depleted by long-term care costs than financial wealth is. This creates an incentive for people to hold more of their wealth in housing than they otherwise would, distorting housing decisions and decreasing the efficiency of the housing market in allocating houses to those who value them most. In particular, there is a disincentive for older people to downsize, even if their homes are very big, because doing so would expose more of their wealth to the risk of being depleted by the LTC costs. As older people with big houses are less likely to sell, prices increase in the housing market and younger people find it more difficult to afford family homes. As such, while the homestead exemption is designed to protect families from having to sell their homes to pay for care it may have important negative effects on the wider population.

To the best of my knowledge, this paper is the first to examine the link between these two policy problems. The main contribution of this paper therefore is in its quantitative analysis of how LTC policy affects housing affordability through an overlapping generations model of the housing market where households face LTC risks. While other papers have offered important insights on the distortions (and protections) provided by the homestead exemption in different settings (Achou 2023; Chang et al. 2023; McGee

2021), they generally estimate life-cycle models for individual households and treat the house price as exogenous, rather than allowing households to interact through the housing market as they pass through their life cycle, thus generating endogenous house prices. As such, the literature to date has not considered what the wider effects of these LTC policies would be on housing market efficiency which could significantly alter the evaluation of their welfare consequences. My key finding, that repealing the homestead exemption brings significant welfare improvements on average equivalent to an increase in annual consumption of £144 (2012 GBP), or 0.6%, fills this gap in the literature and provides important insight into the potential drawbacks of homestead exemption policies.

The paper proceeds as follows. I first set out some descriptive facts regarding housing consumption and moving behaviour over the life cycle in the UK, showing that the ratio of non-durable to housing consumption decreases significantly as households age. I then discuss the policy context around the homestead exemption around the world and the details of the UK setting.

With this established, I then set out an overlapping generations model of the housing market. Households in the model face income, longevity, health and household size risk and make consumption and housing choices every period. They trade a fixed stock of housing, with bigger houses offering more housing services and also allowing agents to adjust the composition of their wealth portfolio towards more housing and where the house price is determined endogenously by aggregate demand and supply. Agents exhibit both temporary and persistent heterogeneity in their preferences for housing.

I estimate this model using the Method of Simulated Moments, matching UK data moments on moving rates and housing choices over the life cycle. I then use the estimated model to evaluate a counterfactual steady state where the homestead exemption is removed and taxes are adjusted to balance the budget. In this counterfactual steady state, house prices are 6.5% smaller. Agents on average receive an increase in welfare equivalent to a £144, or 0.6%, increase in consumption per annum. The biggest benefits accrue to those who had less housing wealth, whereas those with the biggest losses are those with LTC problems or with higher inherited wealth in the initial steady state.

This paper does not model the transition between steady states, and therefore does not permit a full analysis of the likely welfare benefits and costs of a reform to the homestead exemption which also allows for transition costs. These would be particularly relevant in the case of a reform which tends to reduce house prices because of the immediate effect on homeowners with a mortgage, who might face foreclosure. In addition, this paper does not model the supply side of the housing market or try to understand why house supply is apparently so insensitive to demand, rather treating supply as fixed. For these reasons, this paper does not claim to offer a comprehensive analysis of the effects of long-term care policy on housing market efficiency, but rather offers some important first estimates of the extent to which long-term care policy can reduce welfare through distorting the housing market and reducing efficiency.

In doing so, this paper contributes to two distinct literatures. The first concerns saving and consumption decisions in old age especially as they relate to portfolio choices and the effects of the LTC system. Many papers have highlighted the importance of long-term care and medical expenses in driving high rates of saving by older people (De Nardi, French, and Jones 2010, 2016; De Nardi, French, Jones, and McGee 2025; Nakajima et al. 2025). Lockwood (2018) shows the importance of bequest motives in driving low rates of dissaving of wealth while Nakajima et al. (2020), Blundell et al. (2016) and McGee (2021) point to the importance of housing wealth in driving high wealth holdings at older ages. This paper builds on that literature by discussing how the saving decisions of older people, particularly with regard to housing, can have important effects on younger people, and quantifies the welfare cost of policies that distort older people's saving decisions.

The second literature to which this paper contributes is the literature on distortions and efficiency in the housing market (Best et al. 2018; Gervais 2002; Sommer et al. 2018). In particular, this paper is in the spirit of work by Cho et al. (2025), Kaas et al. (2021) and Han et al. (2023) who build quantitative models to assess how costly transaction taxes on housing are. This paper contributes to this literature by analysing a novel type of distortion, namely incentives to save in housing through the LTC system.

The rest of this paper proceeds as follows. In Section 3.2, I discuss the policy context around homestead exemptions in LTC systems. In Section 3.3, I set out an overlapping-generations model of the housing market. In Section 3.4, I discuss estimation and model fit. In Section 3.5, I carry out welfare and counterfactual analysis. Section 3.6 concludes.

3.2 Empirical facts

3.2.1 Data

I use two main data sources throughout this paper, for both the descriptive analysis and the estimation of the model: Understanding Society (UndSoc) and the English Longitudinal Study of Ageing (ELSA). UndSoc is a representative annual panel survey of UK households whereas ELSA is a representative biennial panel survey of the community-dwelling over-50s in England. Both surveys ask respondents a battery of questions regarding housing, income and household choices, with ELSA in particular featuring many questions about LTC needs and beliefs about the LTC system. I use UndSoc Waves 1 to 9 (2009-2018) and ELSA Waves 4 to 9 (2008-2019), to capture a period in between the financial crash of 2007-08 and the COVID pandemic of 2020 onwards, which might lead to abnormalities in the data. In Appendix 3.A.1 I present descriptives for the two samples.

3.2.2 Policy context

3.2.2.1 Housing demand and non-durable consumption over the life-cycle

Figure 3.1 plots log equivalised food consumption (as a proxy for non-durable consumption) and log equivalised housing demand over the life-cycle, using data from UndSoc.

Figure 3.1: Consumption types over the life cycle

(a) Log equiv. food consumption (b) Log equiv. housing consumption

Notes: data from Understanding Society (2009-2018), weighted using household-level weights.

Here, log equivalised food consumption is the log of real food spending[1] divided by the square root of household size, whereas log equivalised housing demand is the log of the number of rooms in the household's home divided by the square root of household size.

There are two notably different patterns here. Consumption is hump-shaped, with the highest consumption in middle age. In contrast, housing demand increases up to middle age and then seems to flatline. In other words, there is no evidence of households adjusting their consumption bundle towards non-durables and away from housing in old age: an increasing proportion of their consumption bundle is made up of housing.

It could be argued that this is due to an increasing preference for housing as agents get older. However, Figure 3.2 plots the change in log equivalised housing demand, conditional on moving (orange, right hand axis), along with moving rates (blue, left hand axis).

What this figure shows is that conditional on moving, older households make very large adjustments to their housing demand, with households who move at age 90 reducing their equivalised housing demand by approximately 34.9% whenever they move. However, they move very rarely at these ages (2.0% of households move house per annum at age 90, compared to 14.0% at age 30) so these changes come very infrequently. While the evidence from this figure is by no means decisive, it is difficult to conclude that pref-

[1] Nominal food spending is deflated to 2012 GBP using the UK CPI (Federal Reserve Bank of St Louis 2025).

Figure 3.2: Movement rates and change in log equivalised housing demand conditional on moving over the life cycle

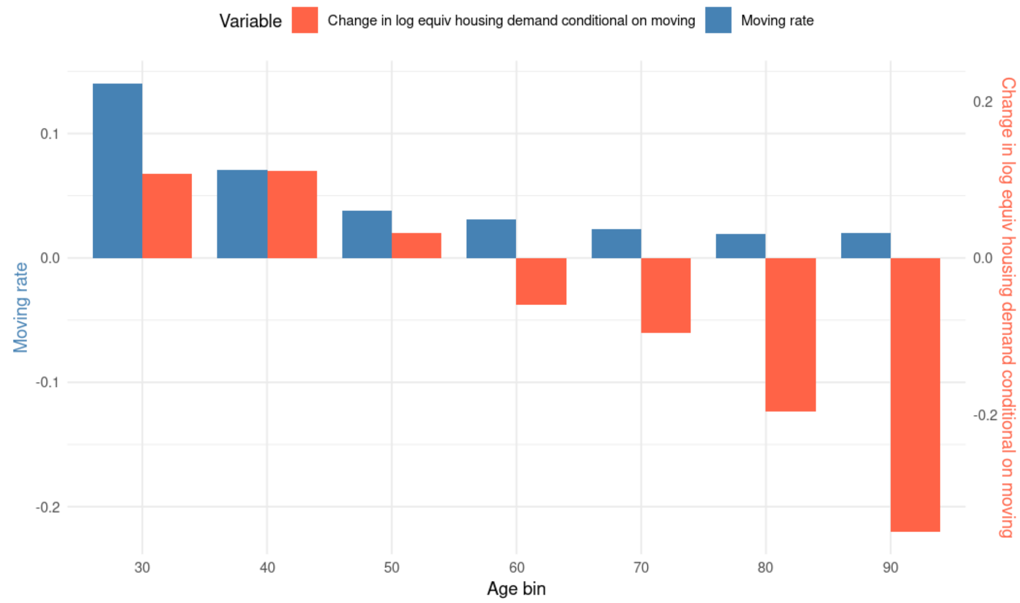

Notes: data from Understanding Society (2009-2018), weighted using household-level weights.

erences for housing relative to non-durable consumption are much higher at older ages, because otherwise households which did move would not make such large downward adjustments to their housing consumption. We might expect, therefore, that insofar as agents' preferences over housing and non-durables can be approximated as being constant over the life-cycle, policies which prevent older people from downsizing encourage them to over-consume housing, which subsequently will drive up the cost of housing for younger people.

3.2.2.2 LTC policy and housing wealth around the world

Many different countries around the world offer means-tested government support with LTC needs. In many cases, housing wealth is treated more generously or exempted entirely from the means test, meaning that it is better protected against LTC costs.

Table 3.1 summarises some key examples of countries which treat housing wealth differently as part of the means test for support with LTC costs. I class these countries together as offering some form of a "homestead exemption", i.e. some form of preferential treatment for housing in the means test.

[2]Note that in the UK the different constituent countries (England, Scotland, Wales, Northern Ireland) have slightly different LTC funding policies, such as different capital limits in the means test for Wales and Scotland relative to England. Where the constituent countries have an aspect of policy in common, such as excluding the primary residence from the means test for support with LTC, I will refer to this as the UK policy, but in cases where the policies diverge I will be specific as to which constituent country's policy I am describing (generally England's).

Table 3.1: Treatment of housing wealth in long-term care (LTC) costs

Country	Treatment of Housing Wealth	Source
Australia	Primary residence not included for aged care means assessment	Services Australia (2025)
France	Primary residence is not counted for eligibility for the Allocation Personalisée Autonomie	Robertson et al. (2014)
Ireland	Only 7.5% of the value of the home is contributed up to a maximum of three years	HSE (2025)
UK[2]	Primary residence is not counted as an asset when determining means-tested support with residential LTC costs	NHS (2025)
US	Primary residence is not counted as an asset when determining means-tested Medicaid eligibility	DHHS (2005)

3.2.2.3 LTC policy and housing wealth in the UK

In the UK system, the means test consists of an asset test and an income test.

In England, the income test allows for agents to contribute their income towards the cost of their care, provided that their residual income is above their Personal Expenses Allowance (PEA, for those in care homes) or above their Minimum Income Guarantee (MIG, for those receiving care in other settings). Certain forms of income are disregarded, notably including income from employment but not income from pensions. In 2025, the PEA was £1.6k per annum and the baseline MIG was £13.7k (£9.0k) per annum for a single person (member of a couple) above the age of 66, with certain adjustments for level of disability or household size. In other words, agents are required to contribute some portion of their (non-disregarded) income towards their care, regardless of their assets. More details are given in NHS (2025).

The asset test, the main focus of this paper, is summarised in Table 3.2. In words, an agent pays out of their chargeable assets until their chargeable assets hit a threshold of £23.25k, after which they pay a portion of their LTC costs until their chargeable assets reach £14.25k, after which the state (in the form of the local authority) covers everything, apart from what they can afford out of their income (i.e. what leaves their income above the PEA or MIG).

A crucial feature of the asset test, noted in Table 3.1, is that the value of the primary residence is included in chargeable assets only if they are receiving care in a care home on a permanent basis and they do not have a spouse or dependant living in their primary residence. In other words, provided the agent or a close relative is staying in the home while the agent receives care, the value of the primary residence is disregarded. The criteria that a dependant or close relative would have to meet in order for this disregard

Table 3.2: Means-tested support based on asset level

Chargeable asset level	Means-tested support
Above £23.25k ("upper capital limit")	No means-tested support.
Between £14.25k and £23.25k	Full support, apart from paying what can be afforded out of income, and paying an extra "tariff income" of £1 per week for every £250 above the lower capital limit.
Below £14.25k ("lower capital limit")	Full support, apart from paying what can be afforded out of income.

Notes: in a couple, each member is generally allocated an equal share of chargeable assets held in common. Figures from DHSC (2025).

to apply, as well as other details of the asset test, are set out in NHS (2025)[3].

This creates an incentive for people to save more in housing wealth than in financial wealth: financial wealth does not enjoy the same protections as primary housing wealth so can be more easily decumulated by LTC costs.

To show how many people are affected by these considerations, Figure 3.3 below plots percentiles of the financial (i.e. non-housing) wealth and total wealth distributions for those aged between 70 and 75, taken from ELSA. The horizontal black line represents the upper capital limit in the asset means test. The graph shows that only around 40% of individuals in this age range have financial wealth above the upper capital limit, and thus would (at least initially) receive no means-tested support with their long-term care costs. However, the 20th percentile of the total wealth distribution is above the upper capital limit, so at least 80% (to be exact, 83%) of individuals would receive no means tested support if there were no distinction between housing and financial wealth in the asset test.

The degree to which the homestead exemption distorts agents' behaviour will obviously be conditioned by agents' knowledge of the long-term care system: if agents are completely myopic about long-term care payment arrangements until they develop care needs then agents' moving decisions at younger ages will not be affected by the incentives of the long-term care system[4]. While there are not questions in ELSA that allow a direct examination of knowledge of the homestead exemption, there are some questions that let us understand agents' beliefs about the long-term care system more generally. Of those

[3]Even if neither the agent nor the close relative are staying in the home, the agent can still defer selling the home to pay for care by applying for a Deferred Payment Agreement, essentially a home equity loan from the government which allows for the agent to pay for care out of their housing wealth without selling their home in their lifetime. The home is only sold after the agent dies to repay the loan (unless the agent or a third party has repaid the loan in the meantime).

[4]Note, however, that even in this completely myopic case agents with care needs would not be forced to sell their home when care costs arise which would clearly impact the supply of housing and thus the housing decisions of younger people.

Figure 3.3: Distribution of financial and total wealth, ages 70-75

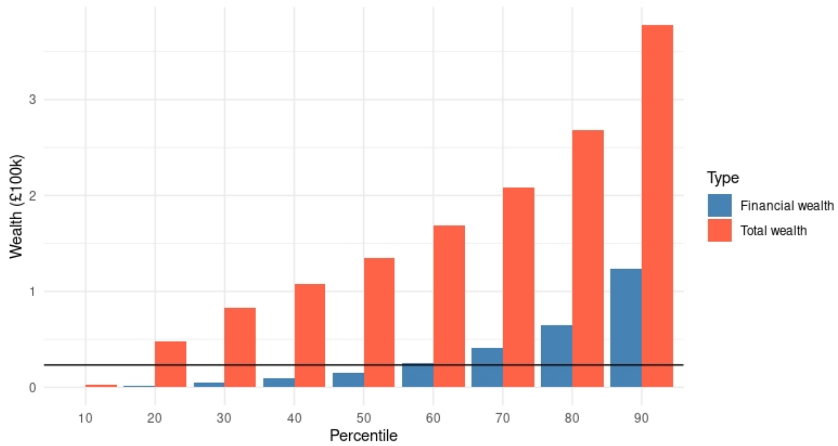

Notes: wealth given at the individual level, with each member of a household being allocated equal share of household wealth. Horizontal black line represents upper capital limit of £23.25k. Monetary values in 2012 GBP. Data from ELSA (2008-2019), weighted using person-level weights.

aged between 50 and 70, 46% say that they have thought about how to pay for their future care costs "a little" (37%) or "in great detail" (9%). 65% of people in this age range say that their own savings will be a source of funding for their future care needs, while 44% say that their local authority will be a source of funding. As such, while forward planning about long-term care costs is by no means universal, there is a large proportion of the population who look ahead to how to fund their social care costs and many of these plan to rely on local authority funding for at least some of their care.

Furthermore, one might object that households are unlikely to change their moving behaviour because of financial incentives of this scale. However, while it is difficult to establish directly a causal link between LTC policy and behaviour because of the lack of variation in the policy over time (NHS 2025), there is evidence that households' moving behaviour is sensitive to more general financial incentives. In particular, Best et al. (2018) and McGee (2021) separately study the impact of the Stamp Duty Land Tax (SDLT), a transaction tax for housing in the UK, on moving behaviour, finding significant sensitivity to changes in transaction costs of buying a house through stamp duty holidays or around notches in the SDLT schedule. For instance, McGee (2021) shows, using an RDD design and ELSA data, that a £5k increase in stamp duty burden causes a 2.85pp reduction, or approximately 50% reduction, in biennial moving rates for the over-50s. As such, while there is not enough variation in policy instruments in the UK setting to study the impact of LTC funding arrangements on moving behaviour in a reduced-form setting, there is reason to believe that the types of financial incentives against downsizing offered by the LTC system could have an important impact on older people's moving behaviour and thus the housing market as a whole.

3.3 Model

3.3.1 Model overview

The decisionmaker in the model is the household. A household consists of either a single person or a married couple. The household can change in size over time as household members die and kids are born and leave the household.

Time is discrete, with each period lasting 5 years. At the start of every period, productivity, health and utility shocks are realised, households receive their income and pay their long-term care costs. Then, each household makes a continuous consumption choice and a discrete housing choice over house size and whether to own or rent.

Multiple generations of households are alive in the model at the same time. There is a fixed stock of housing which is traded in each period, with the price being determined by aggregate demand. When households die, their wealth is passed on to one of the new households in the youngest generation as a bequest.

3.3.2 Resources

3.3.2.1 Liquid wealth

A household earns income through the labour market and - post-retirement at age 65 - through a fixed state-provided pension. Their labour market income is a function of productivity shocks and demographics, including education[5]. They pay taxes on their labour market income to fund government expenditures on pensions, the long-term care system and welfare payments.

Households save in a risk-free asset a_t, which accrues real interest R every period. As such, households' liquid wealth budget constraint will be:

$$a_{t+1} = (a_t + y_t - c_t - \chi_t - r_t + m_t)(1 + R) \tag{3.1}$$

where y_t is income net of income taxes, c_t is consumption, χ_t are long-term care costs, r_t is rent (if the household chooses to rent in the current period) and m_t are the net proceeds from any housing transactions that the household chooses to carry out in the current period.

3.3.2.2 Housing wealth

Agents also can hold wealth in housing. Every period, they choose whether to own or rent, and whether to live in a small, medium or big house. To keep the model tractable, I impose that agents cannot rent medium or big houses, so therefore they can only make

[5] For more details on the income process, see Appendix 3.B.4.

one of four discrete housing choices every period: renting a small house, owning a small house, owning a medium house or owning a big house.

Let $\tilde{h}_t$ represent the number of units of housing stock that an agent owns. Let p_t be the cost per unit of housing. Small houses consist of 1 unit of housing stock, medium houses consist of 1.5 units and big houses consist of 2 units, therefore (for instance) the price of buying a big house is fixed at double the price of buying a small house. The agent's housing wealth is therefore $p_t\tilde{h}_t$.

Households are able to borrow up to some fraction ζ of their housing wealth, which represents their ability to take out mortgages to buy homes. Thus, ζ represents the maximum loan-to-value ratio: if an agent holds housing wealth $p_t\tilde{h}_t$ in period t, the lower bound on a_{t+1} in Equation 3.1 is $-\zeta p_t\tilde{h}_t$[6].

When households purchase a house, they have to pay SDLT on their purchase, with the tax revenue going to the government. They also suffer a fixed utility cost of moving house detailed below.

I assume that if a household does not change their housing consumption between periods, they have not changed house. This means, for instance, that if a household is renting in t and $t+1$ then they are assumed to be renting the same house. This simplifies the model by allowing me to collapse the moving decision and the housing consumption decision into one decision rather than having agents optimise over both.

3.3.2.3 Long-term care costs

Every period, every adult household member - i.e. excluding any children - is in one of three health states: Healthy, Sick or Dead. This implies that the household as a whole is in one of six health states, corresponding to the six combinations (without regard for the order) of the health states of the at most two adults of the household[7]. I estimate transition probabilities between these states as a function of age from the data.

If an individual is sick, they face long-term care costs χ_t. I assume that long-term care needs in this setting imply residential care, i.e. care in an institution, rather than care at home. Individuals can receive government support with their long-term care costs. The rule for calculating government support (when the homestead exemption is in place) is as follows:

- Calculate the agent's share of eligible household assets[8]. Call this sum w_t.

[6]Note that $\tilde{h}_t$ here refers to the housing which the household chooses in period t, not the housing they start the period with - otherwise a renting household would never be able to take out a mortgage.

[7]In other words, the household is in one of the following six health states: Both adults healthy, one healthy and one sick, one healthy and the other dead/absent, both sick, one sick and the other dead/absent, or both dead. In principle, a richer model could track who exactly in the household is sick, which may be relevant if the two adults have different costs of being sick or different probabilities of being sick, for instance by gender. However, I abstract away from these differences for the sake of simplifying the model.

[8]If the agent is in a couple, half of eligible assets are allocated to them. Otherwise, all of eligible

- The agent faces a LTC cost equal to $\min(\chi_t, \max(0, w_t - \kappa))$, where κ is the lower capital limit in the asset means test. In other words, agents are forced to spend down their eligible assets until the lower capital limit[9].

In a world with a homestead exemption, net housing wealth (housing net of mortgage debt) is not included as an eligible asset provided that the house is still being used as a home, e.g. if the agent still has a spouse living in the house while they receive their care in an institution. In this case, eligible assets consist of liquid wealth (excluding any mortgage debt)[10]. If there is no homestead exemption, or if the agent does not have a healthy spouse living in the house, then net housing wealth is included as an eligible asset.

For instance, if the asset test threshold were κ=£20k and an agent has a healthy spouse and their share of household assets is £25k of liquid assets and half of a £250k house with no mortgage, and faces a bill of £100k for their care, then they would pay a sum of £5k=£25k-κ. If there were no homestead exemption (or if they did not have a healthy spouse) they would face the full £100k cost.

If, after paying long-term care costs, a renting household's residual wealth is less than the government-established consumption floor $\bar{c}$, they exhaust their residual wealth and the government tops up their consumption to level $\bar{c}$, capturing means-tested support through the benefit system.

3.3.2.4 Demographics

Households change size over the life cycle as kids are born and leave the household and as (adult) household members die. I do not model couple formation or divorce because of the difficulties involved with keeping track of different households coming together and breaking apart in an overlapping-generations model. Instead, households which start the model single remain single and households which are couples become singles through a household member dying.

To model the presence of children as parsimoniously as possible, I assume there are three possible states with respect to children: either no child has been born to the household, or there is a child currently in the household, or there have been children in the

assets are allocated to them. Note that I calculate household assets after income has been received for the period but before any moving decisions have taken place.

[9] In the real English system, as set out in Table 3.2, there is both an upper capital limit and a lower capital limit, and an intermediate status between these two limits. Also, below the lower capital limit, people still need to contribute out of their income even though their assets are protected. I abstract away from these specific institutional features, keeping only instead a lower capital limit, in order to focus on the key issue of the effects of privileging housing assets in the asset test.

[10] I do not keep track of mortgage debt separately from non-mortgage liquid wealth, instead summarising both in state variable a_t. For the sake of the asset test, I assume that anyone with $a_t < 0$ has mortgage debt equal to a_t and 0 non-mortgage liquid wealth. Anyone with $a_t \geq 0$ has mortgage debt equal to 0 and non-mortgage liquid wealth equal to a_t. As such, agents cannot hold both positive balances of non-mortgage liquid wealth and positive mortgage debt at the beginning of any period.

household in the past but they subsequently left home, i.e. the household has adult children living away from home. The reason for including this third state in the model is that plausibly one reason older households do not downsize in retirement is that they have enough space in their household to welcome their adult children on occasion. I discuss this possibility in more detail when setting out the utility function below.

The probability of having a child in a given period depends on age and whether the household is a couple household. In particular, I estimate from the data the probabilities of switching to having a child at every age by couple status, and use these as exogenous switching probabilities within the model. I also estimate the probability of switching from having a child at home to having no child at home, again as a function of household age and couple status.

Clearly, these assumptions represent a simplification of households' decisions to have children, and mean that I abstract away from the exact number of children and their ages in the household at any given time. However, modelling such complexities - in particular keeping track of the ages of children - would significantly expand the state space of the model with plausibly little improvement in the accuracy of the analysis, so I choose to keep the model tractable by treating households' decisions over children in this simple way. The main consequence of this simplification is that the model will underestimate households' motives for ascending and descending the housing ladder due to household size changes, given that households in the model face less dramatic household size changes than households in the data. As such, if anything, this simplification will tend to lead the model to understate the inefficiency brought about by the homestead exemption because it will understate the variation in willingness-to-pay for housing due to demographics.

3.3.3 Preferences

Agents value consumption, housing services, and bequests. Let s_t be a vector of all state variables for a household at time t: s_t will include demographics, start-of-period liquid wealth and housing wealth, health and productivity. Let d_t be the household's discrete choice over the four housing options at t, which implies the level of housing services they enjoy h_t[11].

For every period when they are alive, a household's utility function is given by:

$$u_a(c_t, d_t, s_t) = g(c_t, h_t) - \omega \times (d_t = rent) - \psi \times (a_{t+1} < 0) - \phi \times (move_t = 1) + \epsilon_t(d_t) \quad (3.2)$$

where $g(c_t, h_t)$ captures a household's base instantaneous utility over consumption

[11] Note that h_t denotes housing services consumed and $\tilde{h}_t$ denotes units of housing owned. These two quantities come apart only in the case of renters, who consume $h_t = 1$ units of housing services but own $\tilde{h}_t = 0$ units of housing.

and housing, ω captures the disutility of renting, ψ captures the disutility of having a negative liquid wealth balance at the end of a period, ϕ captures the disutility of moving house and $\epsilon_t(d_t)$ is a utility shock comprising of a temporary and persistent shock to utility which depends on the household's discrete choice, discussed in more detail below.

The inclusion of the penalty ψ for having a negative liquid wealth balance may seem arbitrary. In a model without such a penalty, households have no incentive not to borrow against their houses, so tend to take out reverse mortgages on their houses in old age in order to fund higher consumption. In the data, households actually tend to pay off their mortgage and then stay mortgage-free until death[12]. The parameter ψ can be interpreted as representing unmodelled costs of applying for and holding a mortgage: for instance, given there is no house price risk in the model, the parameter could be interpreted as capturing the unmodelled default risk, or it could more simply reflect the non-monetary barriers involved with applying for a mortgage. More broadly, it could be interpreted as the disutility associated with dying without having paid off one's mortgage, and thus passing on a non-fully-paid-off house to one's children.

The function $g(\cdot)$ captures instantaneous utility over consumption and housing. I assume it takes the form:

$$g(c_t, h_t) = \frac{((\frac{c_t}{\xi_c(s_t)})^{\alpha}(\frac{h_t}{\xi_h(s_t)})^{1-\alpha})^{1-\lambda}}{1-\lambda} \tag{3.3}$$

where $\xi_c(s_t)$ and $\xi_h(s_t)$ represent equivalence scales for consumption and housing respectively, as a function of state variables s_t. In other words, the household has CRRA preferences over a composite good constructed from equivalised housing and consumption, with α being the weight on consumption.

When the household dies, the household's utility is given by:

$$u_d(beq_t) = \left(\frac{\gamma_1}{1-\gamma_1}\right)^{\lambda} \frac{(\left(\frac{\gamma_1}{1-\gamma_1}\right)\gamma_0 + beq_t)^{1-\lambda}}{1-\lambda} \tag{3.4}$$

i.e. they experience a warm glow from the level of bequests that they give (De Nardi 2004; Lockwood 2018). The parameter γ_0 controls the extent to which bequests are a luxury good and γ_1 controls the strength of the bequest motive. In particular, as set out by Lockwood (2018), in a one-period problem with perfect certainty where households are deciding how much to consume and how much to bequeath of a single good, households with these preferences would choose to bequeath a proportion γ_1 of their assets above the threshold of γ_0.

[12]Appendix 3.A.2 presents and discusses the results of an alternative estimation of the model where ψ is constrained to be 0, demonstrating that in such a model unrealistically many households borrow against their houses late in life.

As such, the value function for a household in period t will be:

$$V_t(s_t) = \max_{c_t, d_t}\{u_a(c_t, d_t, s_t) + \beta[(1 - \delta_t(s_t))\mathbb{E}(V_{t+1}(s_{t+1}|c_t, d_t, s_t)) + \delta_t(s_t)u_d(beq_t|c_t, d_t, s_t)]\} \tag{3.5}$$

where $\delta_t(s_t)$ is the probability of dying at the end of any period, as a function of state variables s_t.

3.3.3.1 Equivalence scales

The equivalence scale for consumption, $\xi_c(s_t)$, is simply defined as the square root of the number of household members, with children being counted as a full household member.

The equivalence scale for housing, $\xi_h(s_t)$ for those who have no children, or who have children in the household, takes the same form, i.e. the square root of the number of household members.

However, the equivalence scale for housing in the case of households who have children who have left the household is somewhat more complicated. As discussed above, one important reason for why older households would choose not to downsize, even though their children have left home, is because there is a benefit to having a bigger house for one's children to visit and stay in. This benefit does not exist for those who have no children either inside or outside the household.

To reflect this, I allow the equivalence scale for housing for those whose children have left the household to depend on both the current household size, excluding any children who have left the household, and the size the household would take if those children returned to the household. In particular, let $\underline{\xi_h}(s_t)$ be the equivalence scale for the household without any children, and let $\bar{\xi_h}(s_t)$ be the equivalence scale for a household with children still in the household. Then, the equivalence scale for a household whose children have left home is given by:

$$\xi_h^{'}(s_t) = (1 - \eta) \times \underline{\xi_h}(s_t) + \eta \times \bar{\xi_h}(s_t) \tag{3.6}$$

i.e. a weighted average of the two equivalence scales, with weight η to be estimated.

3.3.3.2 Utility shocks for housing choices

Every period, households make a discrete choice over housing. These choices are determined both by the consumption and housing services benefits associated with this discrete choice and by the utility shock $\epsilon_t(d_t)$ associated with discrete choice d_t.

In this case, $\epsilon_t(d_t)$ is the sum of two distinct shocks - a temporary shock $\nu_t(d_t)$ and a persistent shock $\rho_t(d_t)$.

The temporary shock is a familiar iid Type 1 Extreme value preference shock with scale parameter σ_ν. Given this temporary shock, if the value associated with each discrete choice i is v_i (inclusive of the persistent shock associated with i), then the probability of discrete choice i being chosen by the household is given by $\frac{\exp(v_i/\sigma_\nu)}{\sum_j \exp(v_j/\sigma_\nu)}$ (Train 2003). The reason for including this shock in the model is to allow for some iid unobserved heterogeneity in preferences for particular housing choices, otherwise the model would not allow observationally equivalent households in the model to make different choices as they do in the data.

The persistent preference shock, $\rho_t(d_t)$, captures heterogeneity in households' attachment to their current home. Some households derive an additional utility flow from remaining in a particular property or location, while others are indifferent. To model this, I assume that a household may receive a positive utility increment $\bar{\rho}$ if it chooses to remain in its current house ($d_t = d_{t-1}$) and zero otherwise. For households without such attachment, $\rho_t(d_t) = 0$ for all choices. Households that would receive $\bar{\rho}$ when staying are said to have a positive preference shock for their current home.

Importantly, this preference shock is persistent. If a household stays in the same house, their attachment carries over into the next period. For instance, if in period t, a household would receive utility shock $\bar{\rho}$ if they stay in the same house and 0 if they move, and they indeed stay, then in period $t+1$ they will again receive $\bar{\rho}$ if they stay and 0 if they move.

If the household moves, they lose this attachment. After moving, they draw a new attachment status for their new house: with a fixed probability, they have a positive preference shock for their new house, otherwise they remain indifferent. In other words, households cannot tell in advance if they will like their new house/area. Consequently, households with a positive preference shock are less likely to move and tend to remain in their current home for multiple periods to avoid losing out on their per-period positive preference shock, in the same way that agents who like a particular house or area are more likely to stay and continue enjoying that house or area.

This feature of the model introduces persistent heterogeneity in housing preferences and ensures that the model does not underestimate the welfare costs of policies that incentivise moving, like repealing the homestead exemption.

3.3.4 The aggregate economy

3.3.4.1 Overlapping generations and the housing market

Multiple generations are alive at the same time within the model. In particular, households enter the model at age 25 and live to a maximum age of 95, so a maximum of 15 generations (at 5 year intervals) can be alive at the same time.

Different generations interact in two different ways: through bequests and through the housing market. Interaction through bequests is simple. When a household dies, their bequest (i.e. the sum of their liquid and housing wealth) is set aside. Each member of the new generation of households (i.e. those starting the model at age 25) draws without replacement from the set of bequests from the households which die at the start of the period, and takes this bequest as their initial condition for liquid wealth. As such, the descendants of homeowners are more likely to be homeowners themselves simply through mechanical intergenerational transmission of wealth.

As for the housing market, there is a fixed supply S_h of housing. The housing market clears at the price when the sum of households' demand for housing is equal to supply. As such, if e.g. older households have higher demand for housing because of the homestead exemption this will push up the price of housing for all households.

I assume that the price of renting is pinned down by a no arbitrage condition:

$$r = p\frac{R}{1+R} \tag{3.7}$$

where r is the per period cost of renting a house and R is the real interest rate. This is because in the absence of capital gains the value of a house as an asset is assumed to be equal to the present discounted value of the stream of rents for the house.

3.3.4.2 The role of government

The government's expenditure is on pension payments to those 65 and over, means-tested support with LTC costs, and benefit payments to those whose income is below the consumption floor. Their revenue comes from income tax payments levied on households of working age as well as SDLT receipts. The proportional income tax rate τ on working-age income is set at a level that would have balanced the budget in the previous period (and thus, in steady state, will balance the budget in the current period).

3.3.4.3 Steady state

To solve for steady state, I use the following procedure:

1 - Select a parameter vector θ.

2 - Solve the individual household problem using backwards induction where preferences are governed by θ. Treat price and tax as state variables in the model, so that policy functions are obtained for a grid of values for p and τ.

3 - Select a candidate house price p.

4 - Interpolate the policy functions calculated from 2 to find approximate policy functions for price p. Simulate the choices of N generations of households using these

policy functions. Find aggregate demand in the final period with a full set of living generations by summing across households' demands[13].

5 - If the aggregate demand in 4 is not equal to (fixed) aggregate supply of housing, go back to 3 and adjust p accordingly.

In doing so, I find the equilibrium price and budget-balancing tax rate which clears the housing market and is consistent with agents' optimising behaviour.

3.4 Estimation

3.4.1 Parameters estimated outside the model

To reduce computational burden I set a large number of parameters outside the model, using either assumed values from the literature or matching moments in data sources such as UndSoc, ELSA or the ONS Life Tables. The parameters that are set outside the model are summarised in Table 3.3 below, with further discussion given in Appendix 3.B.

Table 3.3: External parameters

Parameter	Value	Source
Health state transition probabilities	-	ELSA, ONS (2025a)
Child state transition probabilities	-	UndSoc
LTC cost per period	£54.6k	Dilnot (2011)
Coefficient of CRRA	2	Author's choice
Discount factor over 5 years (β)	0.88	Author's choice
Income process	-	UndSoc
Real interest rate over 5 years	10.4%	Author's choice
Prob. of receiving persistent pref. shock for housing	0.25	Author's choice
Consumption floor	£38.8k	Author's choice
Supply of housing per household S_h	1.432	UndSoc
SDLT schedule	-	HMRC (2025)
Lower capital limit	£14.25k	DHSC (2025)
Maximum loan-to-value ratio (ζ)	0.78	UK Finance (2026)

Notes: see Appendix 3.B for detail on how these values are chosen or estimated.

[13]Note that the other aggregate variable, the tax rate, follows an adjustment rule so that it is equal to the level that would have balanced the budget the previous period, so by the time of convergence to the steady state it is constant.

3.4.2 Parameters estimated inside the model

In total, I estimate nine parameters inside the model. These are parameters for which there are not well established values in previous literature and which cannot be easily set outside of the model. These parameters are summarised in Table 3.4 below.

Table 3.4: Parameters estimated inside the model

Parameter	Description
α	Consumption share of composite good
γ_0	Curvature of bequest motive
γ_1	Strength of bequest motive
ω	Utility penalty of renting
σ_ν	Scale of temporary pref. shock for housing choice
$\bar{\rho}$	Value of positive persistent pref. shock for current house
ϕ	Utility penalty of moving house
ψ	Utility penalty of having negative liquid wealth
η	Equivalence scale weight

Estimation of these internal parameters takes place by the Method of Simulated Moments. I construct the moments using data from Understanding Society (UndSoc) for 2009-2018 and ELSA from 2008-2019. I set out the moments I construct in the following sub-sections and then motivate this choice of moments in Section 3.4.2.4.

3.4.2.1 Moments from UndSoc

UndSoc is particularly useful for constructing moments over the entire life cycle because, unlike ELSA, it interviews respondents over their full life cycle.

I use UndSoc to generate three different types of moments:

- The proportions of households making each of the four discrete housing choices: renting a small house, owning a small house, owning a medium house or owning a big house.
- The proportion of households moving house in every period, both unconditionally and conditional on not moving the previous period.
- The mean difference in demand for big houses between those who have adult children outside of the household and those who have no children.

I define small houses as those with two or fewer bedrooms, medium houses as those with three bedrooms and big houses with four or more bedrooms. Because my model only allows for renting of small houses, I drop any households who are renting medium or big houses (14% of all observations). I evaluate each of these sets of moments using the UndSoc data for each five-year age bin between 25 and 85, apart from the mean difference in demand for big houses by child type, for which I use only the 50 to 85 age bins for reasons of sample size.

In addition, I generate two supplementary moments. First, I use UndSoc data to regress housing demand for homeowners on controls for age and household size as well as individual fixed effects, and treat the variance of the residuals from this regression as an additional moment to be matched.

Finally, I target the price per unit of housing. I set this to be the median real house purchase price in UndSoc per unit of housing bought, using the fact that those who own small, medium and big houses own 1, 1.5 and 2 units of housing respectively.

3.4.2.2 Moments from ELSA

ELSA is particularly useful for constructing moments relating to wealth because it has far more detailed and consistent questions about wealth holdings than UndSoc.

I use ELSA to generate the 25th, 50th and 75th percentiles of the liquid wealth distributions (net of mortgage debt) for each five-year age bin from 50 to 85. I then target as moments in the estimation the proportion of households in the relevant age bin who are below the relevant liquid wealth threshold - for instance, a target in the estimation would be the 25% of model households aged 50 are below the true observed 25th percentile of liquid wealth for 50-year-olds in the data[14].

I also use ELSA to calculate the proportion of households who hold between 0 and £10k in liquid wealth, for each age bin between 50 and 85, and use this as a further set of moments to be targeted.

3.4.2.3 Objective function

In total, this gives 108 moments to be matched. I then find the parameter vector θ which solves:

$$\hat{\theta} = \text{argmin}(\hat{m}(\theta) - m)'W(\hat{m}(\theta) - m) \tag{3.8}$$

where $\hat{m}(\theta)$ is the vector of simulated moments at the parameter guess θ, m is the corresponding vector of empirical moments taken from the data and W is a weighting matrix.

The simulated moments are calculated by drawing from the distribution of households in the age-25 bin in the UndSoc data and creating a household to start the model at

[14] The reason for doing this, rather than trying to match the 25th percentile of liquid wealth in the model to the 25th percentile in the data directly, is that in this alternative case the objective function can be non-smooth.

age 25 with the same observable exogenous state variables, i.e. starting health, couple status, child status and education[15] as the draw from the data. The household is then subject to exogenous shocks to their health, productivity and utility and make choices over consumption and housing over their life cycle, and I use these choices to construct the simulated moments outlined above.

As discussed above, I do not model couple formation, so that agents who start the model single do not then become part of a couple. This makes the model easier to solve but could cause problems insofar as many people who are single when 25 later do go on to be part of a couple. To alleviate this problem, I class agents as starting the model as a couple if they are observed at 25 or at 30 in the data as a couple, so that those agents in the data who were single at 25 and coupled at 30 are counted as starting the model coupled, to avoid having unrealistically many single households in the simulation.

I set the weighting matrix W to be a diagonal matrix with element (i, i) corresponding to the inverse of the variance of the corresponding sample moment i, estimated via bootstrapping.

3.4.2.4 Identification

Here I briefly sketch out how matching these different moments will help identify parameters of interest.

The consumption share of the composite good will be identified by the proportions of people demanding different house sizes. In particular, the less important is consumption (hence more important is housing) in the composite good, the higher demand will be for bigger houses. Note that it is also important to be targeting the price of housing here because the same level of consumption of housing could be brought about by relatively high preference for housing and a high price or low preference and a low price.

The disutility of renting will be identified by the proportion renting versus owning small houses. As an agent enjoys the same amount of housing services whether owning or renting a small house, the key difference in the incentives to take either of these two choices, above and beyond the protected status of housing wealth in the LTC system, is the fixed disutility of renting.

The overall cost of moving will be identified by moving rates in the data, while the

[15] Note that starting wealth at the beginning of the life is not taken from the data but is rather generated by the model, because inheritances are endogenous to the bequests of the people who died the previous period in the model. Also, I impose that everyone starts the model with a random draw from the stationary distribution of the productivity shock and that everyone starts the model without a positive persistent preference shock for their current house, as neither of these two shocks are directly observable in the data. Also, note that in order to start off the OLG model I must make some initial assumption about the state variables of older generations - therefore, for the first period, I populate all living generations in the model using the same method of drawing from the youngest age bin in the UndSoc data, before simulating the model for many periods to remove the influence of these arbitrary initial conditions

value of the positive persistent preference shock will be identified by moving rates conditional on not moving in the previous period. This is because agents who did not move in the previous period will be disproportionately likely to have experienced a positive preference shock for their current house, so the bigger the difference between the unconditional moving rate and the moving rate conditional on not moving the previous period, the more significant is the role of the persistent preference shock.

The strength of the bequest motive will be identified by the extent of liquid wealth holdings in old age, while the curvature of the bequest motive will be identified by matching the 25th, 50th and 75th percentiles of the liquid wealth distribution separately.

The penalty for having a negative liquid wealth balance will be identified by the proportion of households with liquid wealth balances between £0 and £10k. The key idea here is that introducing a penalty for negative liquid wealth holdings encourages bunching at 0 in liquid wealth holdings.

The equivalence scale weight η will be identified by the difference in demand for big houses between those with children who have left the household and those who have never had children. The more similar these two groups behave in their demand for big houses, the more similar the equivalence scale for those whose children have left the household is to the equivalence scale for those who never had children, hence the closer η must be to 0.

Finally, the scale of the Type 1 Extreme Value temporary preference shock will be identified by the variance in housing demand conditional on age, household size and household fixed effects. The bigger this variance in the data, the bigger the role played by unobserved temporary heterogeneity in preferences in the model.

3.4.3 Results

The results of the estimation are given in Table 3.5.

Some of the parameter estimates bear commenting upon. The share of consumption in the composite good is given as 0.716, suggesting that housing services constitute a relatively small proportion of the composite good. McGee (2021) finds that the consumption share of the composite good is 0.567, allowing a greater share for housing, though in his model there is house price risk so it is already less advantageous to hold housing, explaining why the underlying preference for housing would need to be higher to match moments on house choices by households.

The bequest motive parameter estimates suggest a strong bequest motive but that bequests are a luxury. In particular, a household with these bequest preferences[16], if

[16]For ease of exposition, the numerical examples which follow assume a single-good benchmark, i.e. wealth can either be consumed as a non-durable or bequeathed, and there is full weight on non-durable consumption in the utility function. Otherwise, optimal bequests will not depend only on wealth but also on housing consumption.

Table 3.5: Estimation results

Parameter	Estimate
α - consumption share of comp. good	0.716
	(0.004)
γ_0 - curvature of bequest motive	0.457
	(0.005)
γ_1 - strength of bequest motive	0.626
	(0.015)
ω - utility penalty of renting	0.120
	(0.002)
σ_ν - scale of temporary pref. shock for housing choice	0.377
	(0.004)
η - equivalence scale weight for those with absent children	0.127
	(0.001)
ψ - utility penalty for negative liquid wealth	0.096
	(0.001)
$\bar{\rho}$ - value of positive persistent pref. shock for current house	0.494
	(0.003)
ϕ - utility penalty of moving house	0.636
	(0.010)

Notes: estimation via the Method of Simulated Moments. Standard errors estimated from 20 bootstrap replications.

given £Xk at the start of a five-year period and knowing that they will die with certainty at the end of the period, would bequeath 63% of their residual wealth after consuming £45.7k (£9.1k per year) - i.e. they would bequeath $0.63 \times (X - 45.7)$ if $X > 45.7$, and 0 otherwise. This implies that, for instance, a household with final-period wealth of £100k would consume £65.8k and bequeath £34.2k, or a household with final-period wealth of £200k would consume £102.8k and bequeath £97.2k. This is consistent with strong bequest motives estimated in the literature on consumption and saving at old ages (Lockwood 2018).

The parameter for the disutility of renting is slightly positive, so agents suffer a slight utility penalty from renting. One way to interpret the scale of this parameter is to observe that if an agent were indifferent between renting and owning a small house in the absence of any such penalty, the introduction of this penalty for renting implies that the agent would choose to rent 42% of the time[17]. Note that even without a penalty for renting in the model, certain features of the model might lead it to predict low levels of renting relative to the data. One is that the model does not have a great deal of income and wealth heterogeneity or any (regional) price heterogeneity for reasons of tractability, whereas in the data those on persistently low incomes or in very high price areas may never be able to save enough to buy a house so will be "trapped" as renters. The second is that, again for model tractability, there is no house price risk in the model, which makes owning relative to renting more attractive than it might be in the data. For these reasons, the low utility penalty of renting here would best be interpreted as a lower bound on the true value.

The interpretation of the equivalence scale weight η is that the equivalence scale for households with adult children who have left home is composed of 0.127 of the equivalence scale for households with children in the household and 0.873 of the equivalence scale for households who have never had children. In other words, these households behave much more like households who have never had children than households with children still at home when it comes to equivalence scales for housing.

There is a slight per-period utility penalty for having negative liquid wealth, i.e. holding a mortgage. Including such a penalty in the model is key because otherwise households in the model have an incentive to borrow against their housing wealth in old age, which is not observed in the data. In Appendix 3.A.2, I show the results of an alternative estimation where this parameter is fixed equal to 0, and show that the model fit in terms of liquid wealth holdings is considerably worse.

The final two parameter estimates are most easily interpreted together. If, in the absence of any cost of moving, the agent would be indifferent between their current

[17] This figure is arrived at as follows. Let v_0 be the utility of renting and v_1 be the utility of owning a small house. By assumption, $v_0 = v_1 - \omega$. From the estimation, the scale of the Type-1 EV shock is 0.377, so the probability of renting is $\frac{\exp(-\omega/\sigma_\nu)}{1+\exp(-\omega/\sigma_\nu)} = 0.42$

housing situation and moving, the introduction of a utility penalty of $\phi = 0.636$ for moving means that her probability of staying becomes 84% and her probability of moving becomes 16%. If, on top of this, the agent has received a positive preference shock for her current house of size $\bar{\rho} = 0.494$, her probability of staying becomes 95% and her probability of moving becomes 5%[18].

3.4.4 Model fit

Here I show model fit for each of the sets of moments under the estimated parameters.

Figure 3.4 shows how well the model fits the probability of renting over the life cycle.

Figure 3.4: Model fit - proportion renting

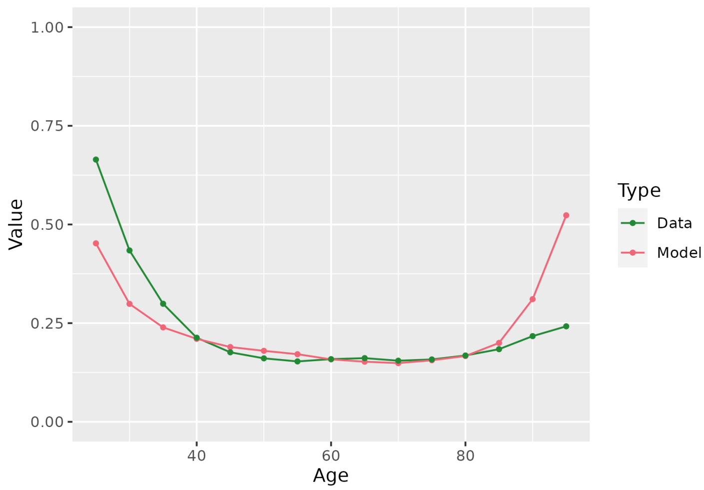

The model matches the level of renting over the middle portion of the life cycle, as well as the feature of the data that renting is higher at the start and the end of life, though the model understates renting propensity at the beginning of the life cycle and overstates it at the end. One explanation for the miss at the start of the life cycle is that there are plausibly many exogenous reasons to move at the start of life, such as moving for careers or study, which are not captured in the model, in which the only reason to move is to change one's level of housing consumption. As for the end of the life cycle, the model overstates slightly agents' propensity to move at old ages (as shown below), which means it overstates the degree to which agents downsize at the very end of life, plausibly because there is no degree of age-dependence in moving costs, whereas in reality it is plausible that older people find it more costly to move than younger people.

Figures 3.5, 3.6 and 3.7 show the proportion of people who own small houses, medium houses and big houses respectively in the model and the data.

[18]These figures are arrived at as follows. Denote the agent's utility from staying as v_0 and from moving as v_1. The probability of staying is then $\frac{\exp((v_0-v_1)/\sigma_\nu)}{1+\exp((v_0-v_1)/\sigma_\nu)}$. If $v_0 - v_1 = 0.636$, this probability is 84% and if $v_0 - v_1 = 0.636 + 0.494 = 1.130$, this probability is 95%.

Figure 3.5: Model fit - proportion owning small houses

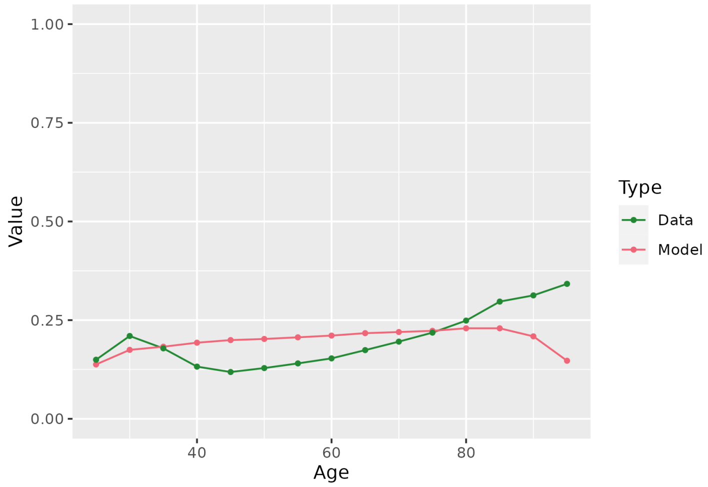

Figure 3.6: Model fit - proportion owning medium houses

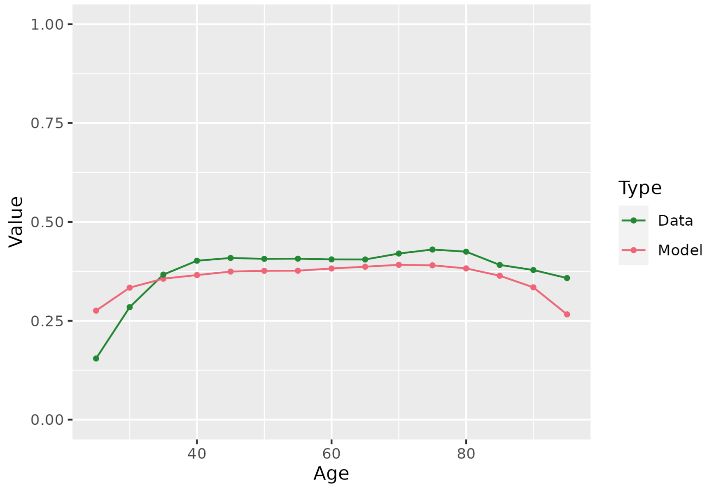

Figure 3.7: Model fit - proportion owning big houses

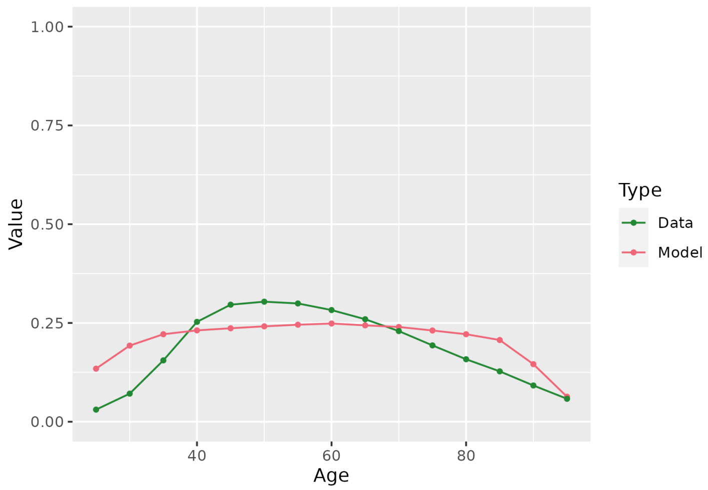

In each case, the fit is reasonably good. The levels of each type of home-ownership in the model is quite close to that of the data, and for medium and big houses the trajectories are broadly similar, insofar as they are both hump-shaped. However, in both these latter two cases, the hump shape is more pronounced in the data than in the model, and the model does not match the inverse hump shape of demand for small houses. One reason for this is that the model is limited in its ability to capture changing household sizes over the life cycle, only allowing a maximum of one child per household, so will not be able to allow for the full demographic pressure to move up and down the housing ladder. The other reason is that the model could be overstating younger households' access to credit and ability to take out mortgages to buy medium and big houses, meaning that the age gradient in demand for these types of houses is steeper in the data than in the model.

Figure 3.8 shows the probability of moving in a given period in both the model and the data. The left-hand panel shows the unconditional probability and the right-hand panel shows the probability conditional on not moving the previous period.

Figure 3.8: Model fit - probability of moving

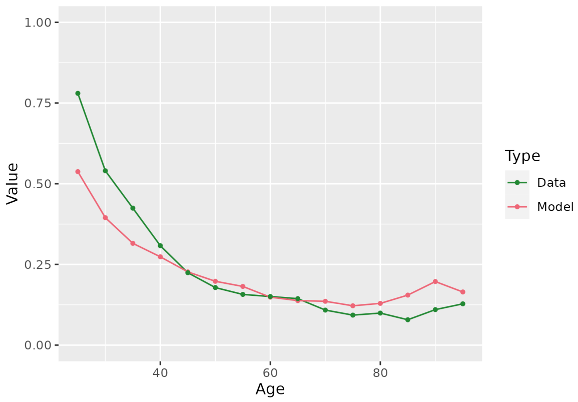

(a) Unconditional

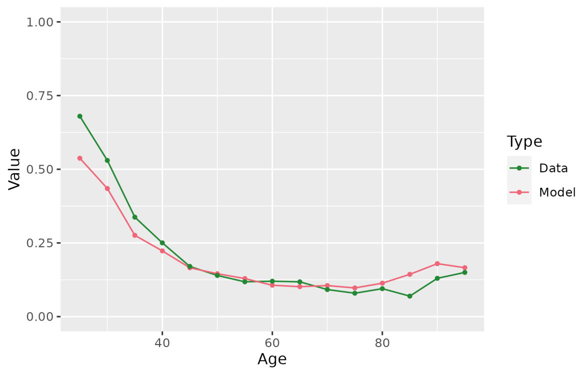

(b) Conditional on not moving in $t-1$

Again, the model is able to broadly match data patterns when it comes to moving. An important caveat here is that agents in the model are only counted as moving if they change their house choice, for instance changing from renting to owning a small house, meaning that if an agent is renting in both $t-1$ and t then they are not counted as moving. In contrast, in the data, agents can move house even when doing rent-rent transitions, which might explain why moving rates in the data (and not the model) are relatively high at young ages before declining. Also, the model slightly overstates moving rates at the end of life, which as mentioned above is plausibly due to a lack of age-dependence in the costs of moving.

Figures 3.9, 3.10 and 3.11 show the model and data results for the 25th, 50th and 75th percentiles of the liquid wealth distribution respectively, with the unit on the vertical axis being £100k.

Figure 3.9: Model fit - 25th percentile of liquid wealth distribution

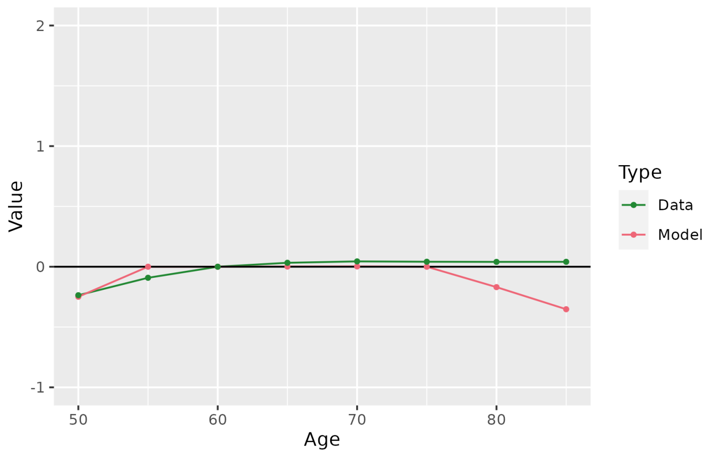

Notes: scale on the vertical axis is £100k. Values shown only from 50 up to 85 age bin because ages in the ELSA data start at 50 and are censored at 90.

Figure 3.10: Model fit - 50th percentile of liquid wealth distribution

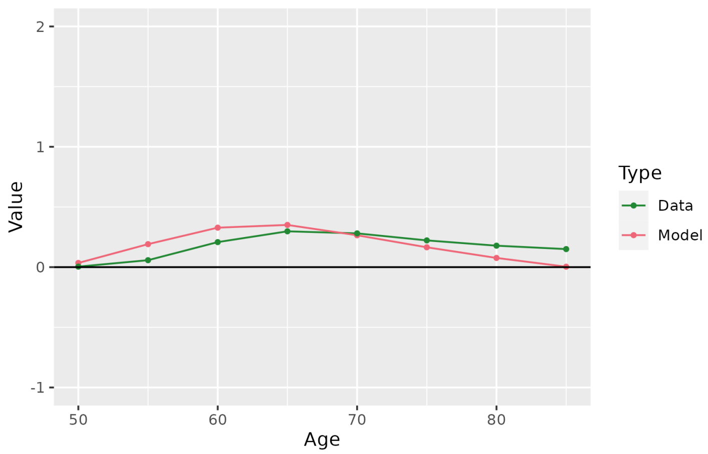

Notes: scale on the vertical axis is £100k. Values shown only from 50 up to 85 age bin because ages in the ELSA data start at 50 and are censored at 90.

Figure 3.11: Model fit - 75th percentile of liquid wealth distribution

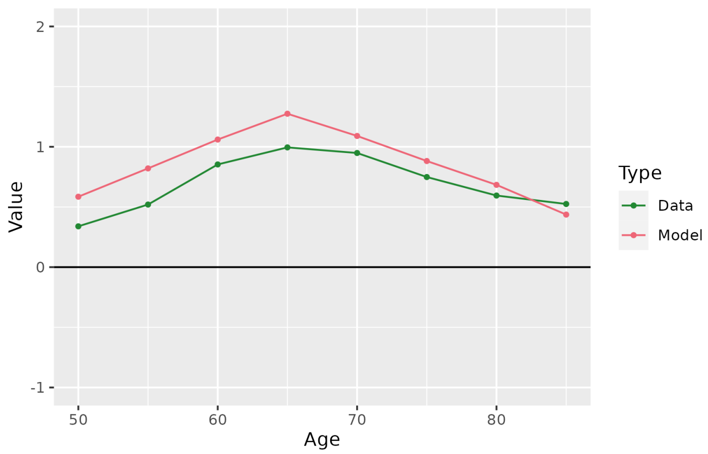

Notes: scale on the vertical axis is £100k. Values shown only from 50 up to 85 age bin because ages in the ELSA data start at 50 and are censored at 90.

For each of the three targeted percentiles, the level and trajectory of liquid wealth is broadly correct. In particular, the inclusion of a penalty for having a negative liquid wealth balance produces bunching at 0 for the 25th percentile of liquid wealth holdings in the model, similar to what is observed in the data. However, even with this penalty, we still observe that the 25th percentile of liquid wealth is negative in the model at older ages, as households borrow against their housing wealth. In Appendix 3.A.2 I show the results of estimation where the penalty for holding negative liquid wealth is set equal to 0, and show that the model fit is considerably worse.

Figure 3.12 shows the model fit with respect to a) the proportion of households bunching around 0 in liquid wealth[19] and b) the difference in demand for big houses between those who have adult children outside the household and those who have never had children.

Figure 3.12: Model fit - bunching around 0 and demand for big houses by child status

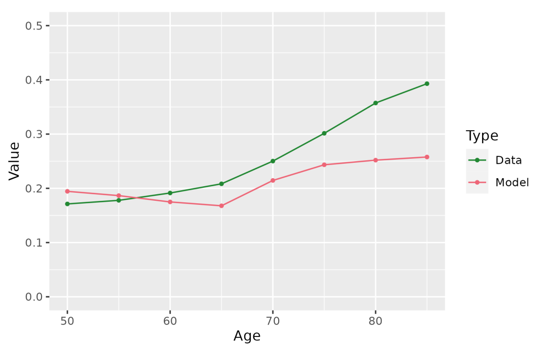

(a) Proportion bunching around 0 in liquid wealth

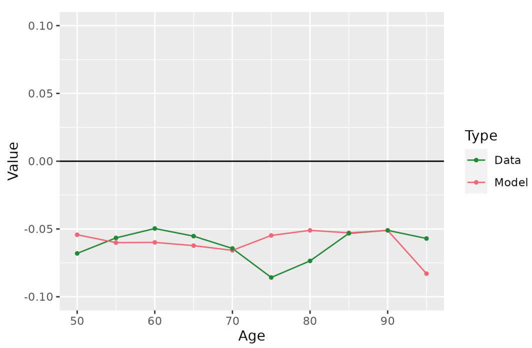

(b) Difference in demand for big houses by child status

[19]In particular, the proportion with between 0 and £10k in liquid wealth.

As for the proportion bunching around 0, the model does not fully capture the degree of bunching nor the full increase in bunching with age. This points to important unmodelled determinants of the decision (not) to borrow against housing wealth in old age which cannot be fully captured with a simple utility penalty for being in debt. The difference in demand for big houses between those with adult children outside of the household and those who have never had children is broadly similar in the model as in the data, however.

The final two targeted moments are the house price and the variance of the residuals from a regression of housing demand on controls for age, household size and individual fixed effects in the data. The house price generated by the model is £156.1k per unit of housing, compared to the target of £153.3k. The data variance in this case is 0.101 and the model variance is 0.155, suggesting that the model is relying on more unobserved heterogeneity than exists in the data to generate these patterns of housing choices.

3.5 Welfare analysis and counterfactuals

The main counterfactual scenario of interest is the case where the homestead exemption is repealed and housing wealth is treated exactly like liquid wealth. To assess this counterfactual scenario, I solve the model under the new policy regime and with a new income tax rate that balances the budget, and find the equilibrium house price in this new steady state.

In this counterfactual scenario, for a given price, aggregate demand for housing has decreased because housing is less favoured as an investment asset. As a result, the equilibrium price must go down to clear the market. Indeed, in the new steady state the house price drops from £156.1 to £145.9k, a 6.5% drop. As now the means test is less generous to households, the budget-balancing income tax that the government needs to levy goes down also from 27.6% to 26.6%[20].

3.5.1 Changes to choices

Figure 3.13 plots house demand over the life cycle in the two steady states with their respective prices. The left hand panel plots the raw means and the right hand panel plots the difference between the means. The units on the vertical axis are units of housing, with a small house being 1 unit, a medium house 1.5 units and a big house 2 units.

The key change is that in the new steady state (blue) agents have higher housing demand at the beginning and end of life and lower housing demand in middle age. This is because in the original equilibrium house prices were high so agents took longer to save

[20] Recall that the income tax needs to fund pension payments and benefit payments as well which is why the change in the tax rate is not proportional to the change in government outlay on long-term care costs.

Figure 3.13: Counterfactual analysis - house demand

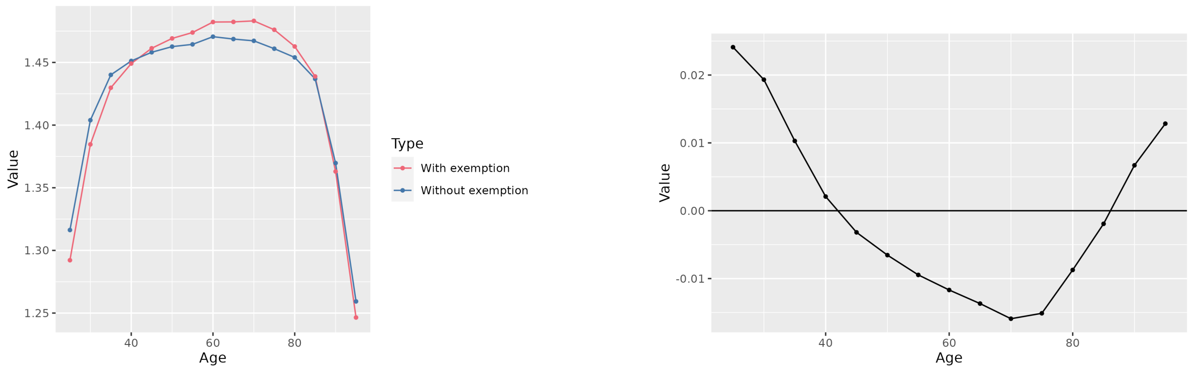

(a) Raw means (b) Difference

up enough to buy bigger houses, and because it was particularly advantageous to hold housing in middle age and early retirement at an age when the risk of long-term care costs was increasing but most households were still couple households and therefore were in a position to benefit from the homestead exemption[21]. The repeal of the homestead exemption removes these incentives and reduces house prices meaning that younger people can climb up the housing ladder faster.

Figure 3.14 shows mean (gross) housing wealth as a proportion of total wealth across the life cycle in both steady states. In this graph, a share above 1 indicates that the household holds a negative liquid wealth position.

Figure 3.14: Counterfactual analysis: housing wealth as a fraction of total wealth

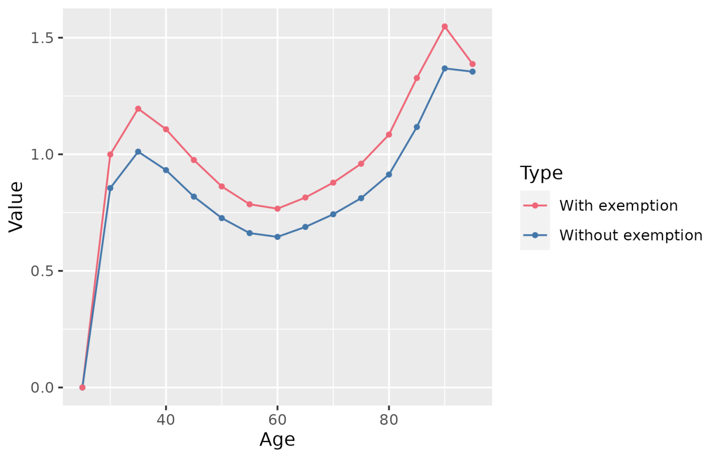

Notes: shares calculated using wealth levels at beginning of period.

In the new (blue) steady state the housing portfolio share is always lower because the

[21]Recall that the homestead exemption only applies if the person with care needs or their spouse is still living at home. At very old ages, most households are single households, and therefore there is no longer an incentive to hold large amounts of housing in the world with the exemption. This, combined with lower prices, explains why at very old ages housing demand is actually higher without the exemption than with the exemption.

extra incentive to hold one's wealth in housing rather than liquid assets has been removed through the repeal of the homestead exemption. Thus, agents adjust their wealth holdings towards liquid wealth and away from housing wealth.

3.5.2 Welfare

I take two approaches to assessing the welfare gains from reform. For the first, to obtain a money-metric measure of welfare gains, I compare average per-period utility in the two steady states, and work out what the average per period increase in consumption would have to be in the old steady state to make the average consumer indifferent between the two steady states, i.e. I find the absolute increase in per-period consumption so that average utility in the two steady states is the same. For the second, to obtain a consumption-metric measure of welfare gains, I find the average *relative* increase in per-period consumption required in order to make the average consumer indifferent between the two steady states.

Overall, repealing the homestead exemption is welfare improving on average. For the absolute, money-metric, welfare measure, the mean welfare gain is £722 per five-year period, or £144 on an annual basis. For the relative, consumption-metric measure, the mean welfare gain is 0.6% of consumption. As a benchmark, mean household consumption in the original steady state is £125k per five-year period, or £25k annually.

There is, however, notable heterogeneity in the gains from the reform. To explore this, I regress the absolute money-metric welfare gain measure I construct through the simulations on a set of regressors capturing the household's demographics in the old steady state in the particular period where I carry out my welfare comparisons. The results of this OLS regression is shown in Figure 3.15. Note that as with any OLS regression with an intercept, a negative value for the coefficient for a particular category does not imply that the welfare gain for people in that category is negative; rather, it implies that the welfare gain for people in that category is lower than for people in the excluded category.

Figure 3.15: Heterogeneity of welfare gains regression

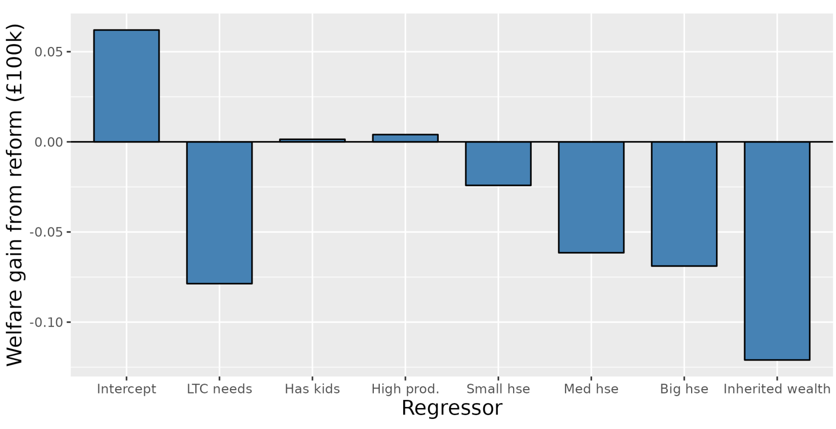

Unsurprisingly, the major losers from the reform are those who are suffering long-term care needs - their wealth is less protected post-reform, so they suffer a significant reduction in welfare. Having kids at home or being in the high productivity state essentially makes no difference to the welfare gains.

Considering the coefficients on owning a small house, medium house and big house, the welfare gain from the reform decreases with the size of the house that a household owns. This is because those with big houses now have a larger proportion of their total wealth left unprotected from long-term care costs, so find the reform less beneficial than other households.

Finally, the coefficient on inherited wealth - i.e. a measure of how much wealth the household starts their life cycle with, through inheritances - is -0.121. In other words, for every extra £1k that the agent started the model with in the pre-reform steady state, their gain from the reform decreases by £121. This suggests that the benefits of the reform tend to accrue to those with less inherited wealth, other things being equal. One reason for this is that the reduced house prices in the new steady state mean that receiving a bequest is less valuable insofar as it is easier to get on the housing ladder without bequests. This points to the repeal of the homestead exemption as potentially reducing the intergenerational transmission of inequality in housing consumption as more households without big inheritances are able to get on the housing ladder sooner.

It should be noted, however, that in the current model, a reform like this does not have any notable effect on the intergenerational transmission of wealth itself because - given the absence in the model of house price growth, and other advantages of housing as an investment asset beyond protection against LTC costs - being able to access housing does not tend to increase one's total wealth.

3.6 Conclusion

In this paper I have shown that the homestead exemption creates important distortions in the housing market by disincentivising older people from downsizing. As a result, house prices (and taxes) are higher and younger people find it difficult to climb the housing ladder, with the housing market being less likely to allocate houses to those who value them most and thus being less efficient. A budget-balanced reform to repeal the homestead exemption reduces prices by 6.5% and increases average welfare by £722 every 5-year period or by £144 per year. There is notable heterogeneity in the gains from the reform, with those with LTC needs, higher inherited wealth and bigger houses in the old steady state gaining relatively less or losing out.

The welfare analysis presented here constitutes a simple comparison of two steady states, abstracting away from the complications of assessing the transition path for the economy between steady states, and thus offers only a partial picture of the full welfare benefits and costs of repealing the homestead exemption. Also, for reasons of tractability the model of this paper simplifies some important features of the housing market. The most important simplification is the lack of house price risk which in other models would alter agents' preferences for holding housing and would expose them to risk of default in their mortgage payments. Moreover, the model of this paper abstracts away from other important sources of risks to households, such as employment risk, while it treats other shocks to the household's utility function (such as the arrival of children) as exogenous whereas more plausibly they would be treated as a decision of the household. Relaxing some of these strong assumptions to allow a more comprehensive treatment of the effects of the homestead exemption on household utility is a promising avenue for future work.

References

Achou, Bertrand (2023). "Housing in Medicaid: Should It Really Change?" In: *American Economic Journal: Economic Policy* 15.1, pp. 1–36. ISSN: 1945-7731. DOI: 10.1257/pol.20200178.

Best, Michael Carlos and Henrik Jacobsen Kleven (2018). "Housing Market Responses to Transaction Taxes: Evidence From Notches and Stimulus in the U.K." In: *The Review of Economic Studies* 85.1, pp. 157–193. ISSN: 0034-6527. DOI: 10.1093/restud/rdx032.

Blundell, Richard et al. (2016). "Comparing Retirement Wealth Trajectories on Both Sides of the Pond". In: *Fiscal Studies* 37.1, pp. 105–130. ISSN: 1475-5890. DOI: 10.1111/j.1475-5890.2016.12086.

Chang, Minsu and Ami Ko (2023). "Marital Transitions, Housing, and Long-Term Care in Old Age". Working paper.

Cho, Yunho, Shuyun Mai Li, and Lawrence Uren (2025). "Stamping Out Stamp Duty: Housing Mismatch and Welfare - The Econometric Society". In: *Quantitative Economics* ().

De Nardi, Mariacristina (2004). "Wealth Inequality and Intergenerational Links". In: *The Review of Economic Studies* 71.3, pp. 743–768. ISSN: 0034-6527. DOI: 10.1111/j.1467-937X.2004.00302.x.

De Nardi, Mariacristina, Eric French, and John B. Jones (2010). "Why Do the Elderly Save? The Role of Medical Expenses". In: *Journal of Political Economy* 118.1. Publisher: The University of Chicago Press, pp. 39–75. ISSN: 0022-3808. DOI: 10.1086/651674.

— (2016). *Savings after Retirement: A Survey.* Rochester, NY. DOI: 10.1146/annurev-economics-080315-015127.

De Nardi, Mariacristina, Eric French, John B. Jones, and Rory McGee (2025). "Why Do Couples and Singles Save during Retirement? Household Heterogeneity and Its Aggregate Implications". In: *Journal of Political Economy* 133.3. Publisher: The University of Chicago Press, pp. 750–792. ISSN: 0022-3808. DOI: 10.1086/733421.

DHHS (2005). *Medicaid Treatment of the Home: Determining Eligibility and Repayment for Long-Term Care*. URL: http://aspe.hhs.gov/reports/medicaid-treatment-home-determining-eligibility-repayment-long-term-care-0 (visited on 11/09/2025).

DHSC (2025). *Social care - charging for care and support 2025 to 2026: local authority circular*. GOV.UK. URL: https://www.gov.uk/government/publications/social-care-charging-for-local-authorities-2025-to-2026/social-care-charging-for-care-and-support-2025-to-2026-local-authority-circular (visited on 11/09/2025).

Dilnot, Andrew (2011). *Fairer Care Funding: The Report of the Commission on Funding of Care and Support*. Tech. rep.

DWP (2025). *Universal Credit statistics, 29 April 2013 to 14 July 2022*. URL: https://www.gov.uk/government/statistics/universal-credit-statistics-29-april-2013-to-14-july-2022/universal-credit-statistics-29-april-2013-to-14-july-2022 (visited on 11/09/2025).

Federal Reserve Bank of St Louis (2025). *Consumer Price Indices (CPIs, HICPs), COICOP 1999: Consumer Price Index: Total for United Kingdom*. URL: https://fred.stlouisfed.org/graph/ (visited on 11/09/2025).

Gervais, Martin (2002). "Housing taxation and capital accumulation". In: *Journal of Monetary Economics* 49.7, pp. 1461–1489. ISSN: 0304-3932. DOI: 10.1016/S0304-3932(02)00172-1.

Gruber, Jonathan and Kathleen McGarry (2025). *Long-Term Care around the World*. Backup Publisher: National Bureau of Economic Research Type: Book. University of Chicago Press.

Han, Lu, L Rachel Ngai, and Kevin D Sheedy (2023). "To Own or to Rent?The Effects of Transaction Taxes on Housing Markets". Working paper.

Hilber, Christian (2023). *How can we make homes more affordable?* British Politics and Policy at LSE —. URL: https://blogs.lse.ac.uk/politicsandpolicy/how-can-we-make-homes-more-affordable/ (visited on 11/09/2025).

HMRC (2025). *Stamp Duty Land Tax rates from 1 December 2003 to 31 March 2025*. URL: https://www.gov.uk/government/publications/rates-and-allowances-stamp-duty-land-tax (visited on 11/09/2025).

HSE (2025). *3-year cap on homes, farms and businesses*. HSE.ie. URL: https://www2.hse.ie/services/schemes-allowances/fair-deal-scheme/3-year-cap/ (visited on 11/09/2025).

Kaas, Leo et al. (2021). "Low Homeownership in Germany—a Quantitative Exploration". In: *Journal of the European Economic Association* 19.1, pp. 128–164. ISSN: 1542-4766. DOI: 10.1093/jeea/jvaa004.

Lockwood, Lee M. (2018). "Incidental Bequests and the Choice to Self-Insure Late-Life Risks". In: *American Economic Review* 108.9, pp. 2513–2550. ISSN: 0002-8282. DOI: 10.1257/aer.20141651.

Lucy Isabella, Makinizi Hoover (2025). *The State of Housing in America*. URL: https://www.uschamber.com/economy/the-state-of-housing-in-america (visited on 11/09/2025).

MACPAC (2026). *Alternative Approaches to Federal Medicaid Financing*. URL: https://www.macpac.gov/wp-content/uploads/2025/04/2025.04-Policy-in-Brief-Alternative-Approaches-to-Federal-Financing-Final.pdf (visited on 03/04/2026).

McGee, Rory (2021). "Old age savings and house price shocks". Working paper.

Mosca, Ilaria et al. (2016). "Sustainability of Long-term Care: Puzzling Tasks Ahead for Policy-Makers". In: *International Journal of Health Policy and Management* 6.4, pp. 195–205. ISSN: 2322-5939. DOI: 10.15171/ijhpm.2016.109.

Nakajima, Makoto and Irina A. Telyukova (2020). "Home Equity in Retirement". In: *International Economic Review* 61.2, pp. 573–616. ISSN: 1468-2354. DOI: 10.1111/iere.12435.

— (2025). "Medical Expenses and Saving in Retirement: The Case of the United States and Sweden". In: *American Economic Journal: Macroeconomics* 17.1, pp. 161–202. ISSN: 1945-7707. DOI: 10.1257/mac.20220211.

National Housing Federation (2025). *Let's fix the housing crisis*. URL: https://www.housing.org.uk/our-work/a-long-term-plan-for-housing/lets-fix-the-housing-crisis/ (visited on 11/09/2025).

NHS (2025). *Financial assessment (means test) for social care - Social care and support guide*. nhs.uk. URL: https://www.nhs.uk/social-care-and-support/help-from-social-services-and-charities/financial-assessment-means-test/ (visited on 11/09/2025).

Nuffield Trust (2025). *What can England learn from the long-term care system in Germany?* URL: https://www.nuffieldtrust.org.uk/research/what-can-england-learn-from-the-long-term-care-system-in-germany (visited on 11/09/2025).

ONS (2025a). *National life tables: UK - Office for National Statistics*. URL: https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/lifeexpectancies/datasets/nationallifetablesunitedkingdomreferencetables (visited on 11/09/2025).

ONS (2025b). *Private rent and house prices, UK: August 2025.* URL: https://www.ons.gov.uk/economy/inflationandpriceindices/bulletins/privaterentandhousepricesuk/august2025 (visited on 11/09/2025).

— (2025c). *Private rental affordability, England, Wales and Northern Ireland - Office for National Statistics.* URL: https://www.ons.gov.uk/peoplepopulationandcommunity/housing/bulletins/privaterentalaffordabilityengland/2024 (visited on 11/09/2025).

Reisenbichler, Alexander (2025). *Germany's housing crisis is making Europe's sick man even sicker.* EUROPP - European Politics and Policy. URL: https://blogs.lse.ac.uk/europpblog/2025/10/31/germanys-housing-crisis-is-making-europes-sick-man-even-sicker/ (visited on 11/09/2025).

Robertson, Ruth, Sarah Gregory, and Joni Jabbal (2014). *The social care and health systems of nine countries*. Tech. rep.

Services Australia (2025). *Aged care means assessment 065-09000000.* URL: https://operational.servicesaustralia.gov.au/public/Pages/older-australians/065-09000000-05.html (visited on 11/09/2025).

Sommer, Kamila and Paul Sullivan (2018). "Implications of US Tax Policy for House Prices, Rents, and Homeownership". In: *American Economic Review* 108.2, pp. 241–274. ISSN: 0002-8282. DOI: 10.1257/aer.20141751.

Train, Kenneth (2003). *Discrete Choice Methods with Simulation.*

UK Finance (2026). *England Mortgage Factsheet.* URL: https://www.ukfinance.org.uk/data-and-research/data/key-mortgage-market-data (visited on 02/09/2026).

Woolley, Frances (2023). "Long-Term Care Financing: What's Fair and Sustainable?" In: *IRPP Study No. 92. Montreal: Institute for Research on Public Policy.*

Yamada, Minoru and Hidenori Arai (2020). "Long-Term Care System in Japan". In: *Annals of Geriatric Medicine and Research* 24.3, pp. 174–180. ISSN: 2508-4798. DOI: 10.4235/agmr.20.0037.

Appendices

3.A Supplementary graphs and tables

3.A.1 UndSoc and ELSA descriptives

For the descriptive analysis in Section 3.2 and the estimation of the model in Section 3.4 I use a combination of Understanding Society (UndSoc) and English Longitudinal Study of Ageing (ELSA) data. I present some descriptive statistics for the two samples below.

Table 3.6: Descriptive statistics

	UndSoc	ELSA
Age (of head)	54	67
Couple status	0.58	0.54
Household members	2.47	1.90
Rooms in house	4.71	4.69
Homeowner	0.62	0.78
N	207335	38945

Notes: Data from 2009-2018 (Waves 1 to 9) of UndSoc and 2008-2019 (Waves 4 to 9) of ELSA, weighted by household-level weights.

3.A.2 Alternative model estimation

The parameter ψ in the model represents a per-period utility penalty for holding negative liquid wealth balances. In this section I re-estimate the model imposing $\psi = 0$ in order to show the important losses in model fit regarding liquid wealth which arise from excluding this penalty from the model. The results of the re-estimation are presented in Table 3.7 below.

The notable changes are that the estimated bequest motive becomes stronger and the preference for housing becomes weaker. This is because in the absence of a penalty for negative liquid wealth balances there must be a different extra incentive to hold liquid wealth wealth, which is achieved by agents in the model saving more for bequests and allocating a greater proportion of that wealth towards liquid wealth.

The model fit is notably worse. In the figures below I compare model and data moments for the targeted percentiles of the liquid wealth distribution at older ages.

Notably, the 25th percentile of liquid wealth never reaches 0 in the model. This is because, absent default risk, there is very little incentive for agents to ever pay off their mortgage[22]. This stands in contrast to the behaviour of agents in the data, who tend to

[22] Agents would want to accumulate liquid wealth for standard precautionary-saving and bequest reasons, but there is no advantage to having a positive versus negative liquid wealth balance in itself.

Table 3.7: Estimation results with no penalty for negative liquid wealth balance

Parameter	Estimate
α - consumption share of comp. good	0.802
γ_0 - curvature of bequest motive	0.360
γ_1 - strength of bequest motive	0.691
ω - utility penalty of renting	0.116
σ_ν - scale of temporary pref. shock for housing choice	0.366
η - equivalence scale weight for those with absent children	0.126
ψ - utility penalty for negative liquid wealth	-
$\bar{\rho}$ - value of positive persistent pref. shock for current house	0.484
ϕ - utility penalty of moving house	0.669

Notes: estimation via the Method of Simulated Moments.

Figure 3.16: Model fit - 25th percentile of liquid wealth distribution

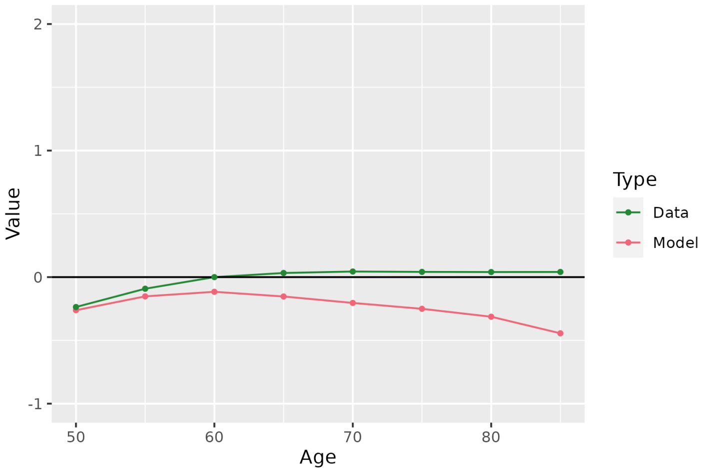

Notes: scale on the vertical axis is £100k. Values shown only from 50 up to 85 age bin because ages in the ELSA data start at 50 and are censored at 90.

Figure 3.17: Model fit - 50th percentile of liquid wealth distribution

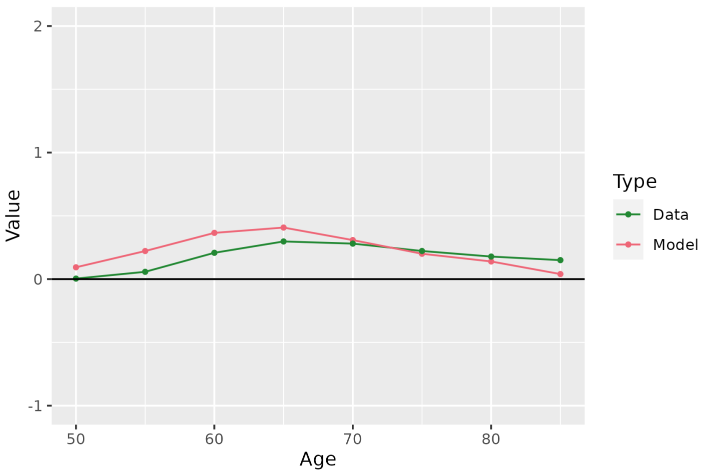

Notes: scale on the vertical axis is £100k. Values shown only from 50 up to 85 age bin because ages in the ELSA data start at 50 and are censored at 90.

Figure 3.18: Model fit - 75th percentile of liquid wealth distribution

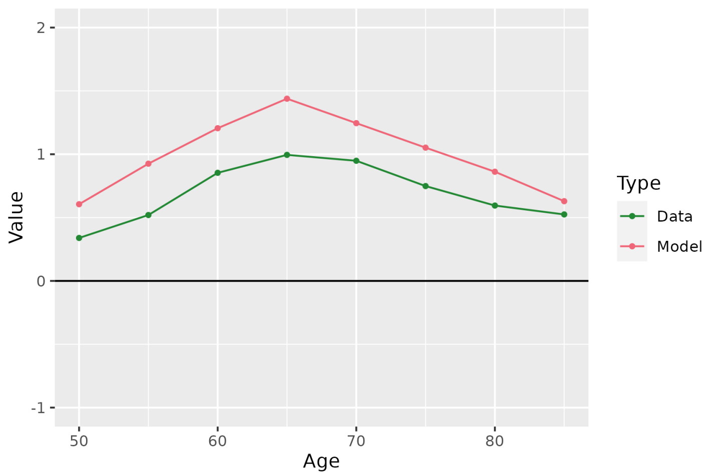

Notes: scale on the vertical axis is £100k. Values shown only from 50 up to 85 age bin because ages in the ELSA data start at 50 and are censored at 90.

pay off their mortgages before retirement.

In the absence of a more complicated model of mortgage choices and default risk, the inclusion of the utility penalty for liquid wealth balances is therefore important for allowing the model to match the data. The penalty itself can be motivated as capturing the unmodelled costs of taking out and holding a mortgage, such as default risk.

3.B Parameters set outside of the model

3.B.1 Health state and child state transitions

I calibrate health state transition probabilities as follows. I use data from ELSA from 2008-2019 to estimate transition rates on an individual basis between being healthy and having LTC needs, where agents are defined as having LTC needs if they are experiencing difficulties with at least two Activities of Daily Living, and are otherwise healthy. I also use ONS life tables (ONS 2025a) to find the probability of dying before one's next birthday at each age for UK adults. I then specify a simple model of health transitions where for each health state i (Healthy, Sick or Dead), the probability of transitioning to health state j at age t is given by $\frac{\exp(\alpha_{ij}+\gamma_{ij}\times age_t)}{\sum_k \exp(\alpha_{ik}+\gamma_{ik}\times age_t)}$, with death as an absorbing state. I estimate these parameters by simulating agents who start healthy at age 25 in order to match the transition rates between states estimated from the ELSA data and the ONS life tables.

For the child state transition probabilities I simply calculate the probability of moving between each of the three child states (Never had children, Child at home, All children have left home), conditional on age and the couple status of the household, for each 5-year age bin in the 2009-2018 UndSoc data, and use these as the probabilities for the model.

3.B.2 Preference parameters

I choose a coefficient of CRRA equal to 2. The 5-year discount factor of 0.88 is derived from an assumed annual discount factor of 0.975. I set the probability of receiving a positive persistent preference shock for housing at 0.25 so that the event of agents developing a strong preference for a particular area is relatively rare.

3.B.3 Care costs

High-quality data on care costs is difficult to find for the UK. Even though ELSA has in more recent waves asked questions about LTC costs, it is a concern that there is likely to be a strong correlation between having high care costs and attriting from the sample.

For this reason, I use approximate measures of care costs from Dilnot (2011), which found that for an adult at age 65 the 20th percentile of future care costs was £0 but the 90th percentile was £100k (in 2009/10 GBP, thus £109.2k in 2012 GBP). By simulating my health transitions model set out above, I find that the agent at the 20th percentile of the distribution of periods spent with care needs after age 65 spends 0 periods with care needs and the agent at the 90th percentile spends 2 periods with care needs. Therefore, I set the cost of one period's worth of support with care needs is £109.2k/2=£54.6k.

3.B.4 Income process

For the income process I use my UndSoc sample to regress the log of household income on a cubic in age (of household head) and dummies for education of household head (college/non-college) and household health states for households below the age of 65, and I set this to be the baseline income for households below 65. For households above 65 (retirees), I set their income to be the mean income in retirement for a household of their couple status, so the income profile over time by couple status is flat.

To capture income risk for those of working age, I assume that households are either in a low- or high-productivity state, where productivity follows a first-order Markov process. Specifically, conditional on being in productivity state i, the probability of continuing in productivity state i is π and the probability of switching is $1 - \pi$. While in the high productivity state, a household's baseline log wage is augmented by μ, otherwise it is decreased by μ.

To calibrate this process I need to find plausible values for π and μ. I do this by simulating households in the model and adjusting π and μ to match i) the persistence in income, as measured by the coefficient on the lagged log income in a regression of log income on lagged income, household fixed effects, number of household members working and a polynomial in age for those households with positive labour supply in both periods in the UndSoc data and ii) the variance of the residual of this regression. The calibration takes place via the Method of Simulated Moments with an identity matrix as a weighting matrix and the result of this calibration is that π is set equal to 0.639 and μ is set equal to 0.149.

3.B.5 Other parameters

I set the consumption floor for households equal to £38.8k, which is the mean Universal Credit payment to households in receipt for 2022, in 2012 GBP, summed across 5 years (DWP 2025).

The SDLT rates are set equal to those as were in place between December 2014 and July 2020, and the lower capital limit is set equal to the level it has been frozen at since 2010, i.e. £14.25k.

I set the real interest rate of 10.4% on a 5 year basis to correspond to a 2% annual real interest rate.

The supply of housing is set to be 1.432 per household. To arrive at this figure, I class houses with 2 or fewer bedrooms as small houses, houses with 3 bedrooms as medium houses and houses with 4 or more bedrooms as big houses. This division of house sizes into three categories is the one which matches most closely my assumption that medium houses provide 1.5 times as much housing services as small houses and big houses 2 times as much, given the distribution of rents by house size in the UK (ONS 2025b). I then calculate the mean housing demand in my UndSoc sample by this measure according to the house sizes demanded in the data and set housing supply equal to this.

The maximum loan-to-value ratio (ζ) is set equal to 0.78, i.e. the mean LTV for first-time housebuyers in England in Q4 of 2025 (UK Finance 2026).

4 | Conclusion

This thesis has presented three stand-alone papers on various themes in the economics of ageing. This conclusion outlines directions for future research building on their findings.

Chapter 1 considered the drivers of the gender care gap in the provision of care by children to parents in the US. The focus of the chapter was primarily on the distributional issue of who provides care and less consideration was given to the complementary question of the efficiency properties of different care arrangements. To make progress in answering this question, it would be necessary to explicitly model the effects of caregiving on the care recipient as well as their preferences for different care arrangements in order to be able to rank care arrangements by agents' preferences. Moreover, the chapter currently models caregiving as a two player dynamic game, and thus is only applicable to a subset of families. Relaxing this restriction would allow more data to be used in estimating the model and would enable the estimation of richer heterogeneity in preferences and costs. While these extensions are outside the scope of Chapter 1 as it stands, and would be difficult to combine with the complex modelling already inherent in Chapter 1, a simpler model of sibling caregiving could allow more detailed consideration of these topics. I leave these as a promising avenue for future research.

Chapter 2 analysed how households respond to mistakes in retirement planning. The key finding of the dynamic model of rational inattention presented there is that inattention about pension ages is ultimately on average not very costly. In the model, costs of mistakes in retirement planning are generated by failures to smooth consumption and because of the limited reversibility of retirement decisions, due to for instance wage and utility penalties associated with rejoining the labour force. However, there are other unmodelled features of household resources or constraints which could determine the costliness of these mistakes: for instance, there may be significant heterogeneity in non-labour income, family support and other forms of insurance, or depreciation of human capital while outside of the labour force, all of which are only partially modelled. Further consideration of these dimensions of heterogeneity - necessarily accompanied, given current computational constraints, with simplifications of other parts of the model - would allow a more complete understanding of the costs of retirement planning.

Chapter 3 assessed the aggregate effect of the "homestead exemption" in the long-term care means test on house prices and welfare. The budget-balanced repeal of the exemption brought decreases in house prices and substantial increases in welfare. A key simplification is that house prices are fixed in steady state, so agents do not face house price or default risk. Other things being equal, this would lead to the model understating true preference for housing because this downside of holding housing is not included. Relatedly, because the model only features a comparison of steady states, it is not able to offer insight on the transitions between steady states and in particular any transition costs associated with a sudden reduction in house prices from the repeal of the exemption. Relaxing these restrictions, and more generally letting more variables in the model to be determined endogenously rather than set exogenously (such as the interest rate and the loan-to-value ratios offered by mortgage providers) would allow a more comprehensive assessment of the effects of the homestead exemption. Computational constraints limited my ability to extend the model in these ways but advances in computational power or simplifications of other aspects of the model could allow further work on these interesting directions for analysis.

With populations ageing across the developed world, the issues considered here will only become more salient in the next decades. This thesis represents a step towards understanding how best we can address the associated policy problems.