Salience and Labour Supply over the Life Cycle

Peter Spittal

Submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

Department of Economics
University College London

June 2020

Declaration

I confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Peter Spittal, June 2020.

Acknowledgements

I am grateful to the many people who have advised or supported me throughout the time I worked on this thesis.

I thank my supervisors, Eric French and Richard Blundell, for their guidance. Eric's time, advice and encouragement in particular have been invaluable. I also thank Jan Eeckhout, Attila Lindner, Fabien Postel-Vinay and Alireza Sepahsalari for their input and support, and Zoe Oldfield and David Sturrock at the Institute for Fiscal Studies for providing access to the data I use in Chapter 3. I acknowledge financial support from the Economic and Social Research Council and the UCL Department of Economics. I thank also my friends and colleagues on the PhD programme, especially Amir Habibi for his valuable friendship and excellent cooking.

I am lucky to have a family who have always supported me. My dad, John Spittal, was as usual a source of calm and constant encouragement. My mum, Wendy Spittal, died shortly after I applied to the PhD programme, but her unwavering and excessive enthusiasm for everything I did still sustains me. I am also grateful to Linda Dawkins, Mel and Pete Bray and the wider Spittal and Bray families for all they have done.

But most of all I thank Lydia, who has felt every extreme of this process every bit as intensely as I have. Her love, support and patience have kept me going. Finally, I dedicate this work to Benjamin, whose arrival in the final stages of its production provided happiness and invaluable perspective.

Abstract

Many policies have dynamic features, linking actions to outcomes at different points in time. Are people aware of these dynamic features, and how does this affect their choices over the life cycle? This thesis contains three chapters which aim to answer these questions.

The first two chapters examine the salience of time limits on entitlement to welfare programmes. In Chapter 1, I identify a feature of policy in the UK which generates a large and foreseeable reduction in benefit income, arising from children ageing out of eligibility for Child Tax Credit. By studying labour supply responses to the benefit reduction, I find evidence that the age-related eligibility rules are non-salient. I also provide evidence of salience increasing through experience. The results also allow me to explicitly rule out a number of competing mechanisms which are indistinguishable from non-salient incentives in other settings but have different implications for welfare and policy.

Then, in Chapter 2, I develop a structural model of life-cycle labour supply in which the age-related eligibility rules may be non-salient. I estimate key parameters of the model by indirect inference, matching the empirical results to identify the proportion of claimants who are uninformed of the rules separately from their labour supply responsiveness. I find that nearly 85 percent of claimants are unaware of the benefit rules, and show that the welfare cost of being unaware that the benefit will expire is substantial.

Finally, in Chapter 3, I study salience of financial incentives in a different dynamic setting: retirement saving. I exploit the substantial heterogeneity in private pension schemes in the UK to study differential salience of the financial incentive to continue working. I estimate peak value models of retirement, allowing for different responses to incentives from different types of pension. I find that similarly-sized labour supply incentives have different effects depending on the part of the system that generates them, indicating differential salience. And labour supply responses are smaller for individuals with more complex pension arrangements (as measured by the number of pensions they hold).

Impact statement

This thesis contributes new evidence that people are unaware of policy features which provide strong inter-temporal incentives for labour supply, savings and consumption. The findings represent academic contributions, as they are among the first evidence that dynamic features of real-world programmes with high financial stakes are imperfectly salient. The evidence is also crucial for effectively designing and assessing the optimality of public policy.

The results in Chapters 1 and 2 contribute new evidence to policy debates—in the UK, and around the world—on the complexity of welfare systems and the potential benefits from simplification. My focus in the chapters is the UK's Child Tax Credit. By studying labour supply behaviour, I show that 85 percent of families entitled to receive benefits under the programme are imperfectly aware that their benefits will fall as their children age out of eligibility, and the resulting failures to plan have costs equivalent to a 14 percent reduction in the life-time value of the benefit.

These findings provide direct evidence of a previously undocumented source of inefficiency in the UK welfare system, arising from imperfect salience of eligibility rules. Recognising the consequent welfare costs substantially reduces the value to claimants. This directly motivates a policy response, as the gains from improving people's awareness of the incentives may be considerable.

The results in Chapter 3 also contribute new understanding to policy debates—in this case, on people's financial preparedness for retirement. I provide evidence that people are imperfectly aware of features of their retirement saving, as manifested in their labour supply responses to financial incentives. This contributes to a growing body of evidence that imperfect awareness in the context of retirement savings. Given the potential implications for living standards in retirement, this also motivates a policy response.

Contents

List of Figures

List of Tables

Introduction

Understanding the way people respond to financial incentives is a central question in public economics. An active literature has documented that, given the complexities inherent in many public policies, people may underweight (or ignore entirely) incentives which would otherwise influence their labour supply, consumption or saving.[1] This has important implications for the design and assessment of optimal policies.

In this thesis, I study the extent to which *inter-temporal* financial incentives arising from policies may not be salient—that is, actively considered by individuals when making decisions. This builds on existing literature which has typically focused on the salience of static incentives. However, the dynamic nature of many policies adds a layer of complexity which may further reduce people's understanding of their incentives. And, as outcomes are typically only realised with some delay in dynamic settings, it may difficult for individuals to learn about the policy features from experience. These features mean there is scope for potentially considerable optimisation errors and, consequently, welfare costs in the context of dynamic policies.

I focus attention on two distinct policies. Both have high financial stakes and provide strong inter-temporal incentives for labour supply, consumption and savings decisions. In Chapters 1 and 2 I study the salience of foreseeable changes to benefit eligibility in the UK welfare system, while the focus of Chapter 3 is retirement saving. My overarching objective across all three chapters is to (i) provide empirical evidence on the extent to which features of the policies I study are non-salient and, in particular, lead individuals to make optimisation errors, and (ii) assess the implications of the findings for welfare and policy design. My approach combines reduced form analysis with structural methods.

In Chapter 1, I provide empirical evidence that the rules linking benefit entitlement to the age of a family's children in the UK are imperfectly salient. These

[1]For example, non-salient incentives have been documented in the context of sales taxes (Chetty et al., 2009; Taubinsky and Rees-Jones, 2017), income taxes (Saez, 2010) and medical insurance (Ketcham et al., 2012).

eligibility rules generate large, foreseeable and lump sum reductions in benefit income. If these rules were salient, the desire of families to smooth marginal utility would lead them to accumulate savings in advance of their child becoming ineligible (by working more and consuming less than they would if the benefits were perpetual), and use these accumulated savings to finance consumption after their benefits fall. In particular, a standard life-cycle model predicts there should be no contemporaneous adjustment to labour supply or consumption at the time they become ineligible (over and above any direct effect of the child's age).

My approach is to learn about the salience of the age-related eligibility rules by testing whether this prediction of no contemporaneous labour supply adjustment holds. I use a difference in differences empirical strategy, designed to identify the effect of reduced benefit entitlement on parents' labour supply over and above any direct effect of their child's age. I find a large and statistically significant increase in labour income due to the benefit reduction—consistent with eligibility rules being non-salient, and families therefore failing to plan for it.

I then provide evidence that the salience of the rules increases through experience, exploiting that families with multiple children experience multiple benefit reductions. Importantly, these auxiliary results allow me to explicitly rule out a number of competing mechanisms which in other settings are indistinguishable from non-salient incentives. For example, models of dynamic inconsistency, frictional labour markets or myopia could not rationalise my results. The ability to distinguish from these alternative mechanisms is a key contribution of the chapter, as the implications for welfare and policy are markedly different.

Then, in Chapter 2, I develop a structural model of life-cycle labour supply to provide further insight into the salience of these age-related eligibility rules. I have two main uses for the model. First, I use the structure of the model to learn more about the nature of the empirical responses I find in Chapter 1. I argue that the structure of the model, combined with the features of my empirical results, allows me to identify (i) preference parameters which govern the size of the labour supply response to a given shock to benefit income separately from (ii) the proportion of households which are unaware of the eligibility rules. Identifying these two effects separately from each other is important for understanding the welfare costs of the non-salient eligibility rules.

I estimate these key parameters of the model by indirect inference, matching the empirical results which provide identifying variation. I find preference parameters which are closely in line with existing literature—reassuring that an otherwise-standard life-cycle labour supply model, with the added feature that rules linking from child's age to benefit entitlement may be non-salient, is capable of generating the empirical results I find in Chapter 1. I also estimate that nearly 85 percent of claimants are unaware of the benefit rules. This underlines that there is widespread non-salience of the rules governing eligibility to a flagship welfare policy with high financial stakes.

Second, I use the model to assess the consequences of the imperfectly salient eligibility rules for the welfare of recipient families. I calculate the compensating variation—the amount that families would be willing to pay, at the beginning of working life, to be informed about the eligibility rules rather than not—as a measure of the costliness of the optimisation errors arising from the non-salient rules over the life cycle. The welfare costs are substantial, equivalent on average to a 14 percent reduction in the generosity of the benefit, even though families learn about the rules through experience. I also show that, in this setting, there are no offsetting benefits to the government from non-salient rules.

Taken together, the analysis in the first two chapters of this thesis provides new evidence that a common feature of welfare policy—linking benefit entitlement to the age of dependents—may be imperfectly salient. Failing to plan for the foreseeable benefit reduction leads families to make significant optimisation errors over their life cycle: they work and save less than they would have liked to if they properly anticipated the benefit reduction. And, while I provide evidence that families do learn about the rules through experience with the system, the resulting welfare costs are still large. My analysis therefore uncovers a previously undocumented source of considerable inefficiency in the welfare system.

In Chapter 3, I shift attention to salience in the context of retirement saving: one of the highest-stakes financial decisions people make over their working life. The rules governing the value of retirement saving schemes, and the consequent incentives they provide for saving and labour supply over the life cycle, are often complex. Given the considerable implications for wealth and living standards in retirement, the extent to which people are well informed about features of their retirement saving schemes is important for policy design. As in the first two chapters of the thesis, my approach is to study labour supply decisions to shed light on the salience of underlying financial incentives—this time, focusing on those provided by pension schemes.

I exploit the heterogeneity in private pension schemes in the UK to estimate the labour supply response to similarly-sized incentives arising from different features of the pension system. If people are equally well informed of all features of their pension portfolio, the labour supply responses should be similar. But if some incentives to continue working are less salient than others, the labour supply responses to the non-salient features would be muted. Differential salience of the incentives would therefore translate into differently sized labour supply responses to similarly-sized financial incentives.

I have three main sets of results. First, I consider the labour supply response to the combined financial incentives provided by all pensions an individual holds. I estimate option value-type models of retirement timing and find a positive and statistically significant impact of marginal incentives to work on continued labour supply. Second, I test for differential labour supply responses to incentives provided by defined benefit (DB) and defined contribution (DC) pensions. Both types of pension are common in the UK, and both can provide similarly sized returns to continued labour supply, but the features of the pensions generating these returns are quite different. I show that labour supply responds more strongly to incentives provided by a DB pension, consistent with the schemes typically being less complex.[2] And so, despite being capable of generating incentives to continue work similar in size to a DB pension, the higher complexity of DC plans may prevent individuals from fully understanding their features. This is a potentially important consideration for policymakers, especially given recent trends towards DC pension plans.

Finally, I provide suggestive evidence that labour supply responses to a given pension incentive are muted for people with more complex pension arrangements. This is consistent with complexity affecting the extent to which individuals are aware of their incentives. Specifically, I use the number of distinct pensions an individual holds as a measure of the complexity of their pension portfolio. There is considerable heterogeneity in the number of private pensions people hold in England: the median member of my sample holds two private pensions, but over 10 percent of the sample holds five or more. I find that the labour supply responses are attenuated for individuals holding more pensions—even when using only within-person variation in incentives for identification. This is consistent with the salience of labour supply incentives provided by pensions decreasing with the complexity of an individual's portfolio.

Overall, the analysis in this thesis provides new evidence of non-salient features of two dynamic, high-stakes, policy relevant environments. It highlights the importance of considering the potential costs arising from individuals misunderstanding (or forgetting about) features of schemes designed to support them, especially when their feedback from the schemes is infrequent—making it difficult to learn from experience.

[2]The value of a DB pension is generally a function only of a measure of late-career earnings, tenure with the scheme, and age relative to a threshold specified in the scheme. By contrast, the value of a DC pension depends on the individual's and employer's contributions, the value of the fund, and investment returns and annuity prices at retirement.

Chapter 1

Benefit Salience: Evidence from the UK

1 Introduction

Tax and benefit systems are complicated. They often result in highly non-linear budget sets, typically through the combination of multiple schedules, thresholds, taper rates and eligibility rules.[1] Given this complexity, people may underweight (or ignore entirely) financial incentives which are not *salient* (e.g. Chetty et al., 2009; Taubinsky and Rees-Jones, 2017). This has important implications for both assessing and designing optimal policy.

In this chapter, I study the salience of *dynamic* incentives provided by the welfare system. I identify a feature of UK policy which leads to a large, exogenous, lump-sum and in principle foreseeable reduction in benefit entitlement, arising from children ageing out of eligibility for Child Tax Credit. I show that the rules governing eligibility are non-salient, despite the high financial stakes, and that claimants learn about the rules through experience. Failing to plan adequately for such a large reduction in benefit income has the potential to generate significant welfare costs; in Chapter 2 of this thesis I show that this is indeed the case. This points to a new source of inefficiency in the welfare system, and highlights the importance of recognising that dynamic features of policy may be non-salient.

My approach is to learn about the salience of the eligibility rules by studying labour supply responses to the benefit reduction. As the benefit reduction is both lump-sum and foreseeable, standard life-cycle labour supply models predict no contemporaneous adjustment in labour supply. The existence (or not) of a causal

[1] For example Mirrlees et al. (2011) discusses the complexity of the tax and benefit system in the United Kingdom, and Joint Committee on Taxation (2015) documents the complexity of the federal tax system in the United States.

labour supply response sheds light on the extent to which claimants anticipated the reduction, and hence whether they were aware of the eligibility rules. The dynamic nature of the setting I study, combined with particular features of the policy, allow me also to investigate directly whether claimants learn about the rules through experience, and to rule out a broad set of alternative mechanisms which are indistinguishable from salience effects in other contexts.

The eligibility rules I study are a feature of a real-world, flagship welfare programme. Child Tax Credit (CTC) was the main source of child-related support for low-income families in the UK for 10 years from its introduction in 2003 until the start of its phased replacement in 2013. Over this period, it reached an average of 5.5 million families each year at an annual cost of £17,500 million (HMRC, 2017).[2] The level of support provided by the policy was substantial: families received £545 a year, plus up to £2,670 a year for each eligible child. This is more generous than other widely studied child-related welfare policies, such as the Earned Income Tax Credit or Child Tax Credit in the US. Children ageing out of eligibility for CTC therefore generated significant reductions in their parents' benefit income.

I use a difference-in-differences empirical strategy to identify the causal effect of reduced eligibility for CTC on labour supply. Using panel data on the labour supply and family characteristics of married mothers, drawn from the UK Household Longitudinal Study (UKHLS) linked to its predecessor the British Household Panel Survey (BHPS), I compare mothers whose benefits are reduced when their child becomes ineligible for CTC against those with children of the same age who experience no change in entitlement.[3]

I first estimate the average labour supply response to a lump-sum reduction in CTC. I find a large and statistically significant increase of £117 a month (standard error £56) in response to a £144 average lump-sum reduction in monthly benefit income. This is inconsistent with the predictions of a standard life-cycle labour supply model, in which households would accumulate savings in advance of a foreseeable lump-sum benefit reduction to smooth consumption and labour supply. But it is consistent with parents having been unaware that part of their benefit entitlement was related to the age of their children and so failing to have anticipated the reduction.

[2]Despite its name, CTC is a transfer paid into claimants' bank accounts each month and is also unrelated to their tax liability. This is unlike the Child Tax Credit in the US, which is administered as a reduction in tax liability and is only partly refundable for families with tax liability lower than their entitlement.

[3]The control group contains mothers whose joint household income or family size mean that their entitlement to CTC is invariant to the child becoming ineligible. I provide evidence that this constitutes a suitable control group below.

Next, I document two further results which are consistent with mothers learning of the eligibility rules through experience, and also rule out a number of alternative explanations for the labour supply response. Both exploit the dynamic nature of the setting I study, along with the fact that mothers with multiple children experience a sequence of similarly sized CTC reductions as each child becomes ineligible.

I show that only mothers who experience a reduction in CTC entitlement for the first time increase labour supply contemporaneously. This is despite that, due to the way CTC is structured, the reduction in benefits is at least as large for subsequent children as for the first child. This suggests that only the first reduction in CTC is unanticipated, with households better able to smooth marginal utility at the time of subsequent reductions. It is therefore consistent with learning about the eligibility rules through experience.

I then estimate separate labour supply effects for families with different numbers of still-eligible children. The *contemporaneous* reduction in CTC does not depend on the number of a family's still-eligible children, but the *future* reductions are higher for those with more remaining children (who will eventually age out of eligibility). This allows for a direct test of whether the labour supply response reflects increased salience of the eligibility rules. If it does, the labour supply response will be larger for families with more still-eligible children, as increased salience of the eligibility rules generates a larger shock to the present value of benefit income for these families. I find that this is indeed the case.

Taken together, these results provide new empirical evidence that a *dynamic* feature of a flagship welfare policy with high financial stakes is imperfectly salient, and that claimants learn about the policy rules through experience. The findings are also inconsistent with a number of other potential mechanisms which could generate a labour supply response to a foreseeable lump sum reduction in benefits, including liquidity constraints, myopia, time inconsistency and misattribution. Being able to distinguish the effects of salience from these alternative explanations, which are indistinguishable in other settings but have different implications for policy and welfare, is a key benefit of this dynamic environment.

However, there are a number of important and interesting questions which cannot be answered directly from the empirical results. Specifically, how widespread is the non-salience of the eligibility rules? Are the average labour supply responses driven by relatively few claimants making a substantial labour supply adjustment, or from smaller adjustment from a higher proportion of claimants? And what are the welfare costs of non-salient eligibility rules in this dynamic setting? I return to these questions in Chapter 2.

This chapter makes two main contributions to the literature on optimisation frictions in tax and benefit systems. First, it provides evidence that a dynamic policy feature, governing entitlement to substantial sums of money as part of real-world welfare policy, is highly non-salient. Existing literature on salience in tax systems has typically focused on static environments which are simpler but with substantially lower financial stakes, such as sales taxes (e.g. Chetty et al., 2009; Feldman et al., 2018; Taubinsky and Rees-Jones, 2017).[4] It is notable to find evidence of a non-salient policy feature in a new setting with high financial stakes, and points to a previously undocumented source of inefficiency in the welfare system.

Second, it provides evidence that people learn through experience. Many studies on non-salient incentives are silent on this question, although there is more general evidence of consumers optimising better over time in relation to Medicare Part D (Ketcham et al., 2012), credit card payments (Agarwal et al., 2008), overdraft fees (Stango and Zinman, 2014), tax returns (Saez, 2010) and cellular bills (Grubb and Osborne, 2015). But, to the best of my knowledge, this is the first study to provide direct evidence that the salience of financial incentives from the tax and transfer system increases following feedback.

This has a number of important implications. It helps distinguish the effects of non-salient eligibility rules from other potential behavioural frictions (such as hyperbolic discounting or myopia) which would not lessen with experience. Being able to identify effects of non-salient rules separately from alternative mechanisms is a key benefit of my setting. The evidence on learning also suggests that the frequency at which people interact with a system is an important determinant of how salient its incentives are: even sizeable financial incentives may be non-salient when they are part of a system people receive feedback from infrequently.

The findings of this chapter also contribute policy debates surrounding the complexity of welfare systems around the world (see e.g. Mirrlees et al. (2011) for the UK). A number of papers have documented that the UK welfare system in general – and the determinants of tax credits in particular – are difficult to understand (McAlpine and Thomas, 2008). This led the UK government to replace CTC and five other existing means-tested benefits and tax credits with

[4] An exception is Dalton et al. (2019), who consider whether observed drug purchases in response to dynamically nonlinear prices as part of Medicare Part D are best explained by a model with non-salient prices or by time inconsistency. By estimating structural dynamic purchase models, they find that a model with non-salient price changes best fit the observed behaviour. However, unlike the environment I study, their setting has no sharp qualitative test to distinguish salience from time inconsistency; their conclusions instead rely on comparing the performance of two structurally estimated models. They are also silent on whether salience effects lessen with experience.

Universal Credit in 2013, motivated by a desire to "simplify the system, making it easier for people to understand" (Department for Work and Pensions, 2015).[5] This chapter provides new evidence that people did systematically misunderstand important features of the welfare system, and this led to a suboptimal labour supply and savings choices over their life-cycle.

Finally, the empirical setting of this chapter is related to a number of papers in labour economics which use predictable changes in child-related benefits to estimate labour supply responses (e.g. Feldman et al., 2016; Wingender and LaLumia, 2017). It is most closely related to Feldman et al. (2016), who also document optimisation frictions in labour supply responses to a similar child-related provision in the US tax code.[6] However, their data do not permit separate analysis of parents who lose eligibility repeatedly, and so their paper is silent on whether people learn about the incentives provided by the benefit system through experience.

The rest of the chapter is structured as follows. I provide further details of CTC in Section 2. I then set out a simple life-cycle labour supply model in which the rules governing eligibility for certain benefits are potentially not salient in Section 3. I use the model to contrast the predictions of the model when the eligibility rules are fully salient against a case where people learn through experience and only become aware of them after experiencing a change in their entitlement. This leads to three testable predictions which distinguish learning through experience from full salience—and a broad set of alternative mechanisms. I discuss the data and empirical strategy for identifying labour supply responses in Section 4. Section 5 sets out my main results and Section 6 contains a number of robustness tests. Section 7 concludes.

2 Details of Child Tax Credit

Child Tax Credit is a means-tested benefit available to parents in the UK living with dependent children. It was introduced in April 2003 alongside Working Tax Credit (WTC) to replace the previous system of financial support provided

[5]Universal Credit also replaced: income-based Jobseeker's Allowance, income-related Employment and Support Allowance, Income Support, Working Tax Credit, and Housing Benefit.

[6]Differently from this chapter, Feldman et al. (2016) document evidence of parents mistakenly treating the increase in their *average* tax rate when their child ages out of eligibility as an increase in their *marginal* tax rate – and so decreasing labour supply immediately after losing eligibility. Differences in the way the US and UK CTCs are administered may account for these different findings. For example, the US CTC is usually administered as a tax refund – and is often claimed in addition to EITC, which may generate different labour supply incentives. By contrast, CTC in the UK is paid as a monthly transfer rather than a reduction in tax liability, and is the primary source of means-tested child related support.

to low-income parents known as Working Families' Tax Credit (Brewer, 2003). Since April 2013, tax credits have been replaced for some potential claimants by the phased roll-out of a new benefit known as Universal Credit.[7] In this chapter, I focus on the 10 year period from April 2003 to April 2013 in which CTC was available to all potential claimants.

Calculation of CTC entitlement. A family's entitlement to CTC is determined by the number of their dependent children (k) and the combined income of co-resident parents (Y). A child is classed as dependent until August 31 after their 16th birthday, or until they child reach age 20 if they remain in full-time non-advanced education.[8] Over the period I study, over 80 percent of children remained in non-advanced education until age 18 – and so, in practice, most families lose eligibility when their child turns 18.[9] When a child no longer satisfies these criteria, their family's k reduces by one.

The maximum amount a family can claim is a linear function of k equal to a constant *family element* (ϕ_f), which all families with at least one dependent child are entitled to, plus a *child element* (ϕ_c) for each of their children. A family's maximum entitlement is therefore given by

$$\bar{b}(k) = \mathbb{1}\,(k > 0)(\phi_f + \phi_c k).^{10} \tag{1.1}$$

In the 2012/13 tax year, the family element ϕ_f was around £545 a year and the child element ϕ_c was £2,690, so a family with two dependent children could claim up to £5,925 a year.

CTC is means-tested against household labour income. For most of the period from 2003 to 2013, a family's total entitlement was withdrawn with household income Y in two stages. First, the child element ϕ_c was withdrawn at rate τ_1 for Y earned above a threshold $\bar{Y}_1$ until it had been withdrawn entirely, leaving only the family element ϕ_f. Second, the family element was withdrawn at rate

[7] Whether a claimant was entitled to CTC or Universal Credit in the period after April 2013 depended on where they lived and a number of other characteristics including whether they had previously claimed CTC and whether they lived with a partner.

[8] Non-advanced education includes study towards academic or vocational qualifications usually provided by schools or colleges (such as GCSEs, A levels, NVQs or BTEC national diplomas). It generally excludes education provided by universities (such as undergraduate degrees) or by an employer as part of a job contract.

[9] In Section 6, I provide evidence that my results are not driven by potential endogeneity in the time a family loses CTC arising from the dependence on when their child leaves non-advanced education.

[10] Following a reform to CTC in April 2017, families are now only able to claim the child element for up to two children. However, for the period I study, a family could claim for all of its children.

Table 1.1: Child Tax Credit Parameters

	Maximum entitlement		Taper thresholds and rates			
	ϕ_f	ϕ_c	$\bar{Y}_1$	τ_1	$\bar{Y}_2$	τ_2
2003/04	£545	£1,445	£13,230	37%	£50,000	6.67%
2004/05	£545	£1,625	£13,480	37%	£50,000	6.67%
2005/06	£545	£1,690	£13,910	37%	£50,000	6.67%
2006/07	£545	£1,765	£14,155	37%	£50,000	6.67%
2007/08	£545	£1,845	£14,495	37%	£50,000	6.67%
2008/09	£545	£2,085	£15,575	39%	£50,000	6.67%
2009/10	£545	£2,235	£16,040	39%	£50,000	6.67%
2010/11	£545	£2,300	£16,190	39%	£50,000	6.67%
2011/12	£545	£2,555	£15,860	41%	£40,000	41%
2012/13	£545	£2,690	£15,860	41%	N/A	N/A

Notes: Table shows the evolution of parameters determining CTC entitlement for eligible families. All figures refer to annual amounts. See the text in Section 2 for details.

τ_2 for income above a second threshold $\bar{Y}_2$. The second threshold was abolished in 2012, after which a family's total CTC claim – including both the child and family elements – was reduced at rate τ_1 until the total entitlement reached zero.[11] Throughout the 10 year period covered by my sample, τ_1 was between 37 and 41 percent and, in all but the year before the second threshold was abolished, τ_2 was 6.67 percent. Table 1.1 shows the parameters determining CTC entitlement in each tax year from 2003/04 until 2012/13, and Figure 1.1 illustrates the CTC schedule for families with one and two dependent children in the 2010/11 tax year.

The taper rate a family faces at given level of income depends on their number of dependent children k through its effect on total CTC entitlement $\bar{b}(k)$: families with more children may face a positive taper rate at higher levels of income than those with fewer children because it takes longer for their total entitlement to be withdrawn. I therefore denote the withdrawal rate a family faces in general as $\tau(k, Y)$. A family's actual entitlement to CTC is given by

$$b(k, Y) = \bar{b}(k)(1 - \tau(k, Y)). \tag{1.2}$$

[11]While I include this change to the way CTC was withdrawn with income in my calculations of CTC entitlement, I do not exploit it directly for identification as the reform happened in the final year of my sample.

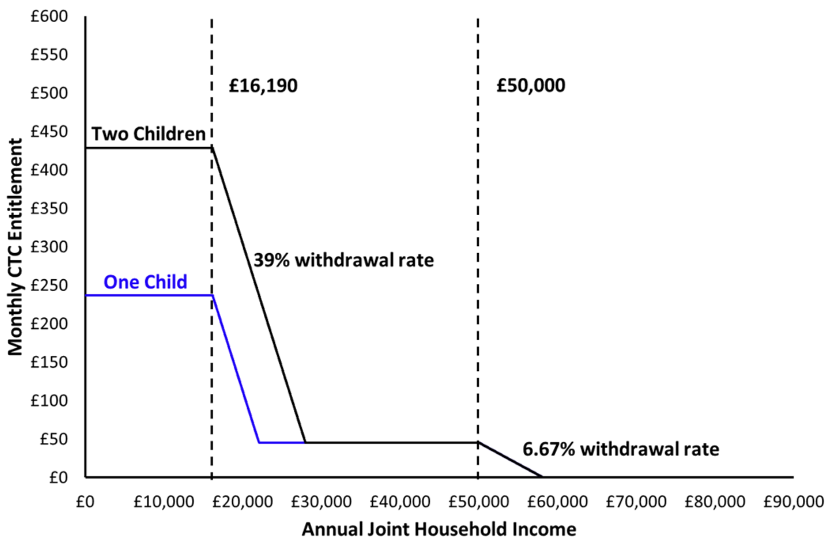

Figure 1.1: CTC schedule in 2010/11 tax year

Effect of a child becoming ineligible. The effect of a reduction in k due to a child becoming ineligible has one of three possible effects on CTC entitlement, depending on joint parental income Y. Some will experience a lump sum reduction in entitlement with no change in the withdrawal rate, others will experience a reduction in both the level of entitlement and the taper rate, while others still will experience no change at all. This variation across parents in the effects of a child becoming ineligible for CTC forms the basis of my strategy to identify the labour supply response to a lump-sum benefit reduction separately from either any substitution effect arising from a change in the withdrawal rate or any direct effect of child age on labour supply.

Comparison with EITC in the US. The tax credits system in the UK plays a similar role to the Earned Income Tax Credit (EITC) in the US by providing means-tested financial support to low-income families. However, a number of key differences make the UK tax credits system particularly well suited to this study. First, the level of child-related support provided by CTC in the UK is especially high, and so a child ageing out of eligibility induces more substantial variation in their parents' benefit receipt than comparable schemes. Over the period 2003 to 2013, families were entitled to claim up to an average of £2,569 for their first eligible child and £2,024 for each of their other children under the UK CTC. By contrast, over the same period the EITC provided up to an average of £1460 for a family's first child, £1120 for the second, and no (or very low) support for any subsequent children.

Second, there was no limit on the number of dependent children a household was able to claim during the period I study. I exploit this feature of the policy directly, as the number of dependent children provides variation in the amount of CTC mothers should expect to lose in future. This allows me to test whether the labour supply response to losing CTC is related to predictable future reductions in eligibility as well as the contemporaneous reduction – a key prediction of the salience of eligibility rules increasing through experience. This would not be possible with a policy such as EITC under which families receive little or no additional payment for their third or subsequent children.

Finally, unlike the EITC, there is no "phase-in" region for CTC, with the maximum benefit amount available to those with the lowest household income. This means that the range of household income for which a family is entitled to the maximum CTC entitlement is wider than for EITC, and these families experience a purely lump sum reduction in benefit income when their child ages out of eligibility. This is important as I am interested in identifying labour supply responses only for this "lump-sum" group.

3 Model

In this section, I outline a simple model in which a unitary household chooses female labour supply, consumption and savings over its lifecycle. A key feature of the model is that the rules linking benefit entitlement to the number and age of a household's children may not be salient, and the salience of the rules may change with experience. I use the model to set out three testable predictions of learning about the eligibility rules of CTC through experience. These predictions form the basis of my empirical analysis.

In Chapter 2 of this thesis, I specify the model fully, with parameterised wage and employment equations and a specific utility function. I then use the empirical results from this Chapter to estimate key parameters within the model and assess the welfare implications of non-salient eligibility rules.

3.1 Overview of the model

Consider a unitary household which chooses how much female labour to supply n_t, and how much to consume (rather than save) c_t in each year t of its working life to maximise expected lifetime utility, subject to its budget constraint.

$$V_t(X_t) = \max_{\{c_j, n_j\}_{j=t,\ldots,\bar{t}}} E_t \Big\{ \sum_{j=t}^{\bar{t}} \beta^{j-t} u(c_j, n_j) | X_t \Big\} \tag{1.3}$$

subject to the budget constraint,

where the expectation is taken over future random variation in risks facing the household given the current state X_t. Supposing that the household can borrow or save freely at interest rate r, its budget constraint can be defined in terms of the asset evolution equation

$$a_{t+1} = (1+r)a_t + w_t n_t + w_t^m n_t^m + T(Y_t, X_t) - c_t, \tag{1.4}$$

$$a_{t+1} \geq \underline{a}_{t+1}, \tag{1.5}$$

with the initial condition $a_0 = 0$ and terminal condition $a_{\bar{t}+1} = 0$, and where a_t denotes the household's assets in year t, w_t is the female wage, $w_t^m n_t^m$ is male earnings and $T(n_t, X_t)$ are net transfers from the government which may depend on household labour supply $Y_t = w_t n_t + w_t^m n_t^m$ and other household characteristics X_t (such as the number and age of its children).[12]

It will be important to specify fully the utility function, wage and employment processes, and sources of risk facing the household for the structural analysis in Chapter 2. However, the only restrictions necessary to deliver the testable predictions which form the basis of my empirical analysis are that within period utility $u(c, n)$ is strictly concave in consumption and female leisure time (defined as residual time spent not working, $1-n$), and lifetime utility is strongly separable over time.

Given these restrictions, both consumption and female labour supply in year t are a function of only the current wage w_t and the marginal utility of lifetime wealth $\lambda(X_t)$ (MaCurdy, 1981). That is,

$$n_t^* = n(w_t, \lambda_t(X_t)), \text{and} \tag{1.6}$$

$$c_t^* = c(w_t, \lambda(X_t)). \tag{1.7}$$

This means, in particular, that female labour supply will not respond contem-

[12]In practice, net transfers may depend on wages and hours worked by men and women in the household separately (not just Y_t). However, this does not affect any of the arguments I make in this section, and so I abstract from it for the sake of notational simplicity.

poraneously to a predictable, lump-sum change in benefit entitlement, because these are already incorporated into the marginal utility of lifetime wealth $\lambda(X_t)$. However, if the reduction in benefit entitlement is unexpected, or otherwise associated with learning that lifetime benefit receipts will be lower than previously anticipated, this will lead to a revision in $\lambda(X_t)$ and a consequent increase in labour supply (and decrease in consumption).

3.2 Salience of child-related eligibility rules.

The household receives transfers (net of taxes) $T(n_t, X_t)$. The salience of the rules mapping from X_t to transfers is the object of study in this chapter: specifically the way the number and age of children determines entitlement to CTC. In particular, I denote the salience of the age-related eligibility rules by $\theta \in \{0, 1\}$. If $\theta = 1$ the household knows how its entitlement depends on the age of its children, but if $\theta = 0$ it incorrectly believes that its entitlement is unrelated.

A household with $\theta = 0$ therefore thinks that its number of children it can claim CTC for in period t will be higher than is actually the case. Non-salient eligibility rules therefore lead a household to over-estimate the level of benefit income it will receive in periods where its actual benefit eligibility will be reduced.

3.3 Testable predictions

The focus of this chapter is mothers' labour supply response to a foreseeable reduction in benefits. I use the model to contrast two scenarios regarding the salience of the eligibility rules: (1) the eligibility rules are *fully salient*, and so $\theta = 1$; and (2) the worker *learns through experience* about the eligibility rules, and so the benefit eligibility rules are non-salient until the worker's first child becomes ineligible, at which point they learn about the way their benefits depends on each of their children's age. In this second scenario, $\theta = 1$ if and only if one of the family's children has aged out of eligibility.

Full salience. The full salience scenario corresponds to a standard life-cycle model with a known sequence of unearned income. In this case, a foreseeable reduction in benefits has no contemporaneous income effect because it is already included in the worker's marginal utility of lifetime wealth. As utility is strictly concave, there would therefore be no labour supply response to a purely lump-sum reduction in CTC.[13] Mothers adjust their labour supply at the time of a

[13] Even if households are fully aware of their future benefit income, labour supply may adjust at the time of a foreseeable benefit *increase* if households are liquidity constrained. However, the optimal plan for a household facing a foreseeable benefit *reduction*, as is the case here, involves

foreseeable change in their benefits only if it affects the benefit withdrawal rate they face and hence their effective wage rate.

Learning through experience. This is not the case if mothers learn through experience. Experiencing a reduction in CTC may lead to contemporaneous income effects in labour supply as parents learn that part of their benefit receipt is related to the age of their children. That is, a reduction in CTC may be associated with a change in θ, which leads to an increase in labour supply through the perceived marginal utility of lifetime wealth $\lambda(X_t)$. When the mother learns about the benefit eligibility rules, she realises that her lifetime unearned income will be lower than she previously expected and increases her labour supply.

Learning through experience has two further testable predictions for mother's labour supply. First, having learned that her benefit receipt depends on the age of her children, there should be no income effect at the time of any subsequent children becoming ineligible. This is because the mother incorporates all future reductions in CTC into her marginal utility of lifetime wealth at the time she learns of the eligibility rules – when any subsequent children become ineligible, the associated reduction in benefits is already incorporated in the mother's labour supply.

Second, the income effect on labour supply at the time of learning about the eligibility rules should be larger for mothers with more eligible children. Mothers with more children yet to become ineligible for CTC experience a larger shock to perceived lifetime wealth when they learn of the benefit eligibility rules, leading to a larger increase in labour supply.

Summary of predictions. There are therefore three predictions of learning through experience which I test directly in the empirical analysis. An increase in labour supply due to an income effect at the time their child becomes ineligible for CTC distinguishes imperfectly salient eligibility rules from full salience. The labour supply response being concentrated at the time of the first CTC reduction, and a larger response for mothers with more still-eligible children, would provide evidence that mothers learn through experience. And, as I discuss in Section 5.4, evidence supporting these latter two predictions also rules out a number of other potential mechanisms consistent with a labour supply response such as myopia, time inconsistency, or misunderstanding the cause of the benefit reduction.

accumulating savings in the period before the benefits fall and then using these savings to finance consumption in the period after. This plan does not involve borrowing, and so any constraints on access to credit do not bind: even a liquidity constrained household would not adjust labour supply at the time of a foreseeable benefit reduction.

4 Data and Empirical Strategy

4.1 Data

I use panel data on mothers' labour supply, personal characteristics and family structure from the UK Household Longitudinal Study (UKHLS) linked to its predecessor the British Household Panel Survey (BHPS), restricting attention to the 10 year period from the introduction of CTC in April 2003 to the start of its phased replacement in April 2013. Interviews are conducted annually, and I use recalled employment histories to construct a dataset recording labour supply and CTC entitlement in each month between interviews. This allows me to study the timing of any labour supply adjustment more precisely than would be possible with annual data.

I calculate the amount of CTC each mother in the sample is entitled to claim and the withdrawal rate she faces based on her household income, hours worked and number of dependent children. I also calculate each mother's counterfactual entitlement and withdrawal rate, i.e. the amount of CTC she would be eligible to claim, and the withdrawal rate she would face, if she had one fewer dependent child. I provide further details on how I constructed the data in Appendix 1.

I focus on the labour supply of mothers because women are commonly found to respond more to changes in tax incentives than men, both in the UK and in other countries (e.g. Blundell and MaCurdy, 1999; Meghir and Phillips, 2010). Failing to anticipate a reduction in CTC is therefore likely to result in a larger labour supply adjustment for women than for men, making it easier to detect. In the main analysis, I further restrict attention to mothers who cohabit with a partner (who I refer to as "married"). This provides an additional source of variation in CTC, which is assessed based on the combined income of both parents – married mothers who earn the same as each other, but have spouses who earn different amounts, may experience different reductions in CTC when their children age out of eligibility.[14] However, I show in Section **??** that the pattern and magnitude of all my results are robust to also including single mothers in the sample.

I make a two further restrictions on the sample. First, I keep only mothers whose children become ineligible for CTC during the period covered by the data – i.e. I exclude mothers whose children are all too old to have ever qualified for CTC, or are all too young to have become ineligible by April 2013. Second, I focus on labour supply in the 24 months either side of each child becoming

[14] In cases where a child lives with one parent and their (non-parent) partner, the cohabiting partner's income still counts in determining CTC entitlement.

Table 1.2: Descriptive Statistics

	All		Lump Sum		MTR		Control	
	Mean	S.D.	Mean	S.D.	Mean	S.D.	Mean	S.D.
Mother's age	47.1	5.7	45.9	5.8	48.2	5.4	48.1	4.8
Child's age	18.82	1.75	18.69	1.79	18.84	1.72	19.19	1.63
Number of children	2.7	1.2	3.1	1.4	2.2	0.9	2.4	0.8
Labour income (£/mo.)	927.8	1049.0	532.4	709.0	1117.0	899.5	1762.6	1539.7
Probability work	0.73	0.44	0.59	0.49	0.86	0.35	0.88	0.33
CTC amount (£/mo.)	179.3	233.5	312.6	258.0	62.3	100.0	16.3	43.5
CTC withdrawal rate	13.18	16.91	17.53	18.51	11.62	15.43	2.82	6.23
Mothers	1204		657		548		227	
Children	1,870		945		635		290	
Obs (children × months)	72,531		36,001		25,183		11,347	

ineligible for CTC.

The final sample includes 1,204 mothers with 1,870 children who lost eligibility for CTC between April 2003 and April 2013. Table 1.2 contains descriptive statistics for the final sample as a whole, as well as for the subsamples I describe in Section 4.2.

4.2 Empirical Strategy

Children becoming ineligible for CTC changes their parents' benefit schedule – reducing the level of benefits they receive and, for those on the CTC withdrawal region, also reducing the effective marginal tax rate. The main objective of this chapter is to test whether this change in benefits causes a contemporaneous labour supply adjustment, even when the change is only lump sum. I also test two further implications of the learning through experience model outlined in Section 2: whether mothers who lose eligibility for CTC multiple times (because they have more than one child who becomes ineligible) adjust their labour supply each time, or whether any responses are concentrated on the first occasion; and whether mothers with more remaining eligible children respond more than those with fewer children, reflecting that mothers with larger families become aware of larger future benefit reductions.

The key empirical challenge is to separate the causal effects of reduced benefit income from any direct effects child ageing may have on maternal labour supply. Parents may decide whether or how much to work based on their family composition, leading to systematic labour supply adjustments at the time their child becomes ineligible for CTC for reasons other than the change in benefits. I propose a difference in differences empirical strategy to identify the effect of reduced benefits on labour supply. I compare mothers who experience a change to their benefit schedule when their child becomes ineligible for CTC against those with children of the same age who experience no change in CTC entitlement.[15] As all mothers in the analysis have children of the same age, this holds constant the direct effect of child age on labour supply.

I divide the sample into mutually exclusive groups based on the effect a child becoming ineligible for CTC has on the household's benefit schedule. I assign parents to groups by comparing their *actual* CTC entitlement in the 24 months before their child becomes ineligible for CTC to the *counterfactual* entitlement they would have faced if their child was already ineligible. This procedure holds household income fixed when comparing the actual and counterfactual entitlements, ensuring that parents are assigned to groups based on the causal effect of child ineligibility on benefit receipt.

There are three groups. Group C is the control group of mothers who experience no change in CTC when their child becomes ineligible. These are mothers whose combined household income puts them on a section of the CTC schedule where their actual and counterfactual entitlements are identical in each of the 24 months before their child becomes ineligible (i.e. they are entitled to only the family element or are entirely ineligible to claim). Requiring that the benefit entitlement of control group parents is invariant to child eligibility in each of the 24 months before the child becomes ineligible is conservative: if a parent's benefit entitlement would be different if their child was ineligible in even one of the 24 months before the child actually becomes ineligible, I consider them to be potentially treated.

There are two main treatment groups. Group L contains mothers who experience only a lump sum reduction in CTC when their child becomes ineligible. These mothers have a counterfactual entitlement below their actual entitlement for at least one of the 24 months before their child becomes ineligible, but have identical actual and counterfactual withdrawal rates in each month. The behaviour of this group is the main focus of the chapter as it sheds light on whether they anticipated the benefit reduction. If these mothers did anticipate the loss of CTC, they should not adjust in their labour supply at the time of the reduction. But, if they failed to anticipate losing entitlement for CTC, they would increase their labour supply due to an income effect. As there is no substitution effect for this group, the existence of a labour supply response is consistent only with the parents having failed to anticipate the loss of benefits.

By contrast, group M experience a reduction in both the level of benefits and

[15] This second group is made up of mothers whose joint household income in the period before losing entitlement to CTC puts them on a part of the CTC schedule which is invariant to one child becoming ineligible, either because their income is too high to receive any CTC, or because they receive only the *family element* described in Section 2.

the effective marginal tax rate when their child becomes ineligible for CTC. This group contains mothers whose actual and counterfactual withdrawal rates differ in at least one of the 24 months before their child becomes ineligible. A labour supply response among this group confounds any potential income effect with a substitution effect induced by a reduction in the withdrawal rate. The strategy of this chapter is to test for the existence of an income effect, for which group L is sufficient; I estimate responses for group M separately only to separate any response of this group from L.

This procedure for constructing groups L and M is deliberately conservative in the sense that the criteria for being included in group L are as stringent as possible. By assigning individuals to group M (rather than L) if their actual and counterfactual withdrawal rates would have differed in any one of the 24 months before losing eligibility, group M is likely to contain some individuals whose actual reduction in CTC was in fact lump sum. This ensures that group L does not contain mothers whose withdrawal rate is affected by their child becoming ineligible, and provides reassurance that the group I study do not experience a reduction in their withdrawal rate.

I estimate labour supply responses for each treatment group separately by estimating equations of the form

$$y_{it} = \alpha + \sum_{g \in \{L,M\}} \beta_g (T_i^g \times A_{it}) + \sum_{g \in \{L,M\}} \gamma_g T_i^g + \delta A_{it} + \zeta X_{it} + \theta_t + \epsilon_{it}, \quad (1.8)$$

where $\{T_i^g\}_{g \in \{L,M\}}$ are binary variables equal to 1 if the mother of child i belongs to treatment group g, and A_{it} is a binary variable equal to 1 if child i is ineligible for CTC in month t. X_{it} are exogenous characteristics of the mother including her age, highest qualification and total number of children, and θ_t is a year-specific fixed effect. The parameters $\{\beta_g\}_{g \in \{L,M\}}$ measure how mothers in each of the two treatment groups adjust their labour supply y_{it} (measured as either monthly labour income or a binary employment variable) differently from those in the control group in the period after their child becomes ineligible for CTC, conditional on characteristics X_{it}. If benefit eligibility rules are fully salient, there should be no labour supply response to a purely lump sum reduction in CTC and so $\beta_L = 0$. But if parents fail to anticipate the reduction in benefits, they may increase their labour supply due to an income effect and so $\beta_L > 0$.

The model of learning through experience set out in Section 2 has two further testable implications. First, there should be an income effect only on the labour supply of mothers who have not had a child previously lose eligibility for CTC.

Second, any income effect should be larger for mothers with more children who are still eligible for CTC. To test for these effects, I split groups L and M into further subgroups: first depending on whether the mother has already had a child lose eligibility for CTC, and then based on the number of the mother's remaining dependant children. I estimate versions of equation (1.8) which allow for separate treatment effects between these subgroups.

Interpreting estimates of β_g as the causal effects of CTC changes for treatment group g requires that (1) the direct effect of a child becoming ineligible for CTC on their mother's labour supply is the same across the control and treatment groups (a parallel trends assumption) and (2) losing eligibility for CTC is unrelated to unobserved determinants of the mother's labour supply conditional on observables X_{it} (an exogeneity assumption). I provide evidence supporting both requirements in Section 6.[16]

5 Results

I establish three main results about the labour supply response of married mothers to a lump sum reduction in benefits induced by their child becoming ineligible for CTC. First, I show that there is a causal effect of losing CTC on maternal labour supply in the months immediately after their child becomes ineligible. Next, I show that all of the income effect is driven by mothers who lose eligibility for CTC for the first time. The estimated income effect at the time any subsequent children become ineligible is small and statistically insignificant. Finally I show that the income effect estimated at the time mothers lose CTC eligibility for the first time is larger for those with more children who are yet to age out of eligibility for CTC.

The results are consistent with each of the three predictions of learning through experience derived from the model in Section 2 and, in particular, are inconsistent with the eligibility rules being fully salient. And, as I discuss in Section 5.4, the results are also inconsistent with a number of alternative models which could potentially generate income effects at the time of a foreseeable lump-sum change in benefits.

[16]The basis of my empirical strategy is that the age-related eligibility rules for CTC provide the exogenous variation in entitlement required by condition (2). However, this is potentially jeopardised for parents of children aged between 16 and 20, who retain eligibility only if their child remains in full time education. If unobservable determinants of these mothers' income also influence whether their child remains in education between ages 16 and 20, this would break the independence between ϵ_{it} and A_{it} in equation (1.8). In Section 6.1 I show that my conclusions are unaffected by using an instrumental variables strategy designed to isolate the exogenous variation in eligibility arising from child ageing, supporting the interpretation of my results as causal effects.

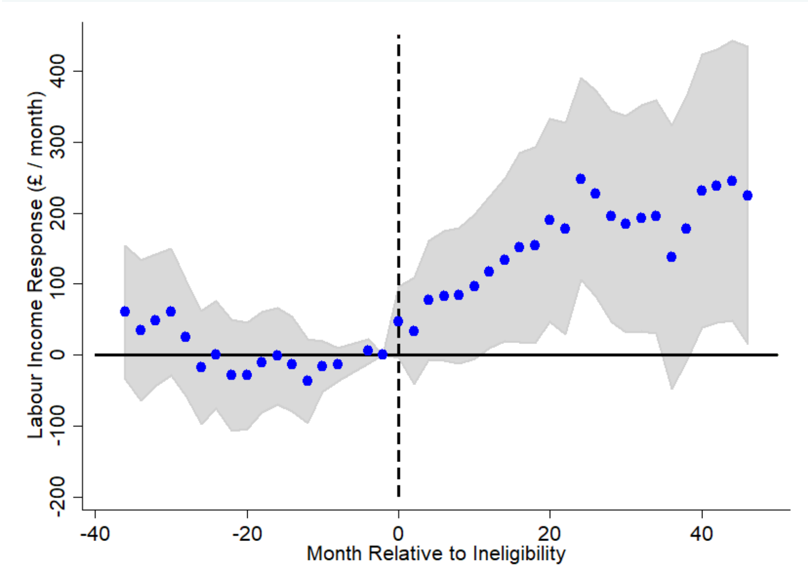

Figure 1.2: Labour Supply Effect of Losing CTC.

Notes: Figure shows the difference in maternal labour income between treatment group L and control group C in each two-month period around their child becoming ineligible for CTC. Each dot is the estimated β_L coefficient from equation (1.8), amended to allow separate effects in each two-month period. The grey area shows the 90 percent confidence interval. The coefficients are normalised relative to the 1-2 month period before the child becomes ineligible.

5.1 Income effect following loss of CTC

I begin by examining the causal effect of a child becoming ineligible for CTC on their mother's labour supply, focusing on those in treatment group L who experience a purely lump sum benefit reduction. As a first look for any potential treatment effect, Figure 1.2 compares the labour supply of mothers in the control group C against those in treatment group L. The labour income of the groups evolves similarly in the months before losing eligibility for CTC, diverges sharply when the children become ineligible for CTC, before levelling off after around two years. This suggests that there is a causal effect of the benefit reduction on the labour supply of mothers in the treatment group, even though the benefit reduction they experience is lump sum and, in principle, foreseeable.

Next I estimate the causal effect of losing eligibility for CTC using regression equation (1.8), again focusing on mothers in treatment group L whose child ageing out of eligibility for CTC leads them to experience a purely lump-sum reduction in benefit income. I present the results in Column (1) of Table 1.3. Panel A contains the estimated treatment effects on mothers' monthly labour income, probability of working, monthly CTC receipt and the withdrawal rate. These are estimates of β_L from equation (1.8), estimated on data from 24 months before the child becomes ineligible for CTC to 24 months after for each dependent variable. All specifications include controls for mother's age, education and total number of children, as well as calendar year fixed effects.

I find a large and statistically significant increase in labour income of £117.3 a month (s.e. £56.1) in response to losing eligibility for CTC. Labour supply responds less on the extensive margin, with a statistically insignificant increase of 2.1 percentage points (s.e. 2.0). Together, these results point to a substantial intensive margin response, with mothers who already work increasing labour income in response to losing CTC. I also estimate the treatment effects of losing eligibility on the CTC schedule. Mothers in group L lose an average of £144 a month in benefit payments, while the change in the withdrawal rate is very small and statistically insignificant. This underlines that, by design, mothers in treatment group L experience only a lump-sum change in their benefit income when their children become ineligible for CTC.

In Panel B, I express these estimated effects as elasticities with respect to unearned income.[17] I distinguish two cases. The "myopic" elasticities assume that all of the estimated labour supply adjustment is solely in response to the contemporaneous change in unearned income. The resulting labour income elasticity is -0.61 and the extensive margin elasticity is -0.09, although the latter is not statistically significant at conventional levels. However, if mothers learn about CTC eligibility rules through experience, these elasticities may overstate the sensitivity of labour supply to changes in unearned income. In that case, the estimated labour supply increases would be a response to both the contemporaneous reduction in CTC and the realisation that benefits may fall further in future.

The "forward-looking" elasticities take this into account. In these calculations, I scale the estimated percentage labour supply adjustment by the total proportion of CTC in unearned income, rather than just the estimated contemporaneous reduction. That is, I assume that mothers adjust their labour supply as if they had lost *all* of their remaining CTC at the time their child became ineligible. These elasticities reflect, in a simple way, that the effect of losing CTC on mothers' perceived lifetime budget may exceed the actual contemporaneous

[17] I have calculated these elasticities by (1) calculating the percent changes in labour supply and unearned income by scaling the treatment effects by the relevant labour supply variable or total unearned income in treatment group L in the 24 months before losing CTC, (2) dividing the percent change in labour supply by the percent change in unearned income, and (3) bootstrapping with 500 samples to estimate standard errors.

Table 1.3: Labour Supply Response to Lump Sum CTC Reduction

	All	First Child	Other Child	First, Large Family	First, Small Family	Other Child
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Treatment Effects (β_L)						
Labour income	117.3**	158.9***	19.57	258.0**	150.5**	36.19
(£/ month)	(56.15)	(58.99)	(60.74)	(107.6)	(60.42)	(57.42)
Probability of	0.0209	0.0377*	-0.0185	0.0573	0.0351*	-0.0138
working	(0.0202)	(0.0212)	(0.0303)	(0.0619)	(0.0213)	(0.0252)
CTC amount	-144.3***	-142.6***	-148.3***	-150.8***	-142.8***	-147.1***
(£/ month)	(1.910)	(1.913)	(2.020)	(3.817)	(1.973)	(2.114)
Taper rate	0.000	0.001	-0.001	0.000	0.001	-0.001
	(0.0004)	(0.0004)	(0.0008)	(0.0014)	(0.0005)	(0.0008)
Observations	72,531	72,531	72,531	72,531	72,531	72,531
Panel B: Implied elasticities						
Labour income	-0.612*	-0.767**	-0.116	-2.459**	-0.597**	-0.435
(Myopic)	(0.334)	(0.332)	(0.382)	(1.135)	(0.277)	(0.450)
Labour income	-0.224*	-0.291**	-0.0404	-0.527**	-0.268**	-0.152
(Forward-looking)	(0.123)	(0.127)	(0.134)	(0.244)	(0.125)	(0.158)
Participation	-0.0893	-0.146*	0.0936	-0.551	-0.0622	-0.0366
(Myopic)	(0.0876)	(0.0843)	(0.159)	(0.465)	(0.0736)	(0.115)
Participation	-0.0327	-0.0553*	0.0325	-0.118	-0.0280	-0.0128
(Forward-looking)	(0.0321)	(0.0321)	(0.0553)	(0.100)	(0.0331)	(0.0400)
Panel C: Income replacement rates						
Myopic	0.813**	1.114***	0.132	1.638**	0.999**	0.495
	(0.386)	(0.403)	(0.416)	(0.687)	(0.392)	(0.450)
Forward-looking	0.297**	0.422***	0.0459	0.351**	0.450**	0.173
	(0.142)	(0.153)	(0.146)	(0.150)	(0.176)	(0.159)

Notes: Panel A contains the estimated treatment effects on various dependent variables, estimated on data from 24 months before the child becomes ineligible for CTC to 24 months after. All specifications include controls for mother's age, education and total number of children, as well as calendar year fixed effects. For further details see the main text in Section 5. Panel B expresses the labour supply treatment effects as elasticities, and Panel C expresses the treatment effect on labour income an income replacement rate. Standard errors, clustered at the level of the mother, in parentheses. Standard errors in Panels B and C calculated by bootstrapping with 500 repetitions.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

reduction.[18] This assumption reduces the labour income elasticity to -0.22 and the extensive margin elasticity to -0.03.

Finally, in Panel C, I provide the income replacement rate implied by the treatment effect estimates. The myopic estimate reflects simply that the £117 a month increase in labour supply offsets 81 percent of the £144 contemporaneous CTC reduction, while the forward-looking estimate shows that the labour supply increase offsets nearly 30 percent of the total amount of CTC mothers will lose when all of their children have aged out of eligibility.

The results presented in this section are contrary to the predictions of a standard lifecycle model, which predicts that there should be no labour income response at the time of a predictable lump-sum reduction in unearned income. However, the findings are consistent with the eligibility rules being initially non-salient, and so mothers failing to anticipate the reduction. In the following subsection I present evidence that, while initially non-salient, mothers learn of the eligibility rules through experience.

5.2 Income effect only when first child loses eligibility

A testable implication of experience with benefit eligibility rules increasing their salience is that any income effect on labour supply should be strongest the first time mothers have a child become ineligible for CTC. If losing eligibility for the first time makes the ways benefit income is related to the age of their children more salient, mothers would adjust labour supply to account for both the contemporaneous and any future expected changes to benefit income arising from eligibility for CTC. Correspondingly, there should therefore be no additional labour supply treatment effect at the time subsequent children age out of eligibility.

I divide each treatment group L and M further into (i) those who experience a reduction in CTC entitlement for the first time, and (ii) those who have already had a child become ineligible for CTC. The result is four new treatment groups $\{L_{first}, L_{other}, M_{first}, M_{other}\}$. I then re-estimate equation (1.8) for these four treatment groups and test whether the labour supply response within group L is different for mothers who lose eligibility for the first time and those who already have experience of the system.

I present the results for group L_{first} in column (2) of Table 1.3, and for L_{other} in column (3). The estimated treatment effect on both labour income and the probability of working is large and statistically significant for mothers who lose eligibility for CTC for the first time: they increase labour income by an average

[18] I also note that, over a sufficiently long horizon, the the delay between learning about the future benefit reduction and the benefits actually being reduced becomes negligible.

of £158.9 a month, (s.e. £59.0) and there is a 3.8 percentage point increase in the probability of working (s.e. 2.1). Both estimates are larger than the effects for all of group L reported in column (1). By contrast, the effects are close to zero and statistically insignificant for mothers who have already had a child become ineligible for CTC. This is despite the reduction in CTC being at least as large for those in group L_{other} as for those in L_{first} because parents lose the constant family element of CTC in addition to the child element when their last child becomes ineligible. These results are consistent with the overall treatment effect reported in column (1) being driven by those parents whose first child ages out of eligibility, with a substantially smaller (and indistinguishable from zero) effect for mothers whose subsequent children become ineligible.

Panel B expresses these estimated treatment effects as elasticities. The myopic labour income elasticity is -0.77 for the first child and the extensive margin elasticity is -0.146; the forward-looking elasticities are mechanically smaller at -0.30 and -0.06 but remain statistically significant. All elasticities are numerically larger than those estimated for all of group L presented in column (1). By contrast, the elasticities for other children in column (3) are numerically small and statistically insignificant, reflecting the statistically insignificant treatment effects for this group.

Finally, Panel C presents the income replacement rates implied by the estimates. The myopic replacement rate, which treats the labour income increase as a response solely to the contemporaneous CTC reduction, is 111 percent for group L_{first}: that is, mothers increase labour income more than the contemporaneous benefit loss. This very high replacement rate points to mothers adjusting their labour supply to more than just the contemporaneous benefit loss, but instead realising that benefits will fall further in future and adjusting labour supply to reflect these future losses as well. Correspondingly, the forward-looking replacement rate is smaller at 42 percent. This means that, at the time their first child becomes ineligible, mothers increase labour income to replace 42 percent of the total amount of CTC they will eventually lose when all of their children become ineligible. The myopic and forward-looking replacement rates estimated at the time subsequent children lost eligibility in column (3) are numerically small and statistically insignificant, underlining that all of a mother's labour supply response occurs when their first child becomes ineligible.

5.3 Income effect larger for mothers with more dependent children

If the salience of benefit eligibility rules increases through experience, and so mothers adjust labour supply to reflect both the contemporaneous benefit income reduction and the realisation that benefit income will fall further in future, the labour supply response should be larger among mothers with the higher future benefit reductions. The CTC setting allows for this implication to be tested directly by comparing the labour supply responses of mothers with different numbers of children still eligible for CTC.

I test whether the labour supply response to a mother's first child becoming ineligible for members of group L is different for those with relatively more dependent children than for those with fewer. I define a family as "large" if it has more dependent children than the median in group L (i.e. three children or or more), while a "small" family has two children or fewer.[19] I then divide treatment group L into three ($L_{first,large}$, $L_{first,small}$ and L_{other}) and allow the labour supply treatment effects to differ across these groups.

Columns (4) to (6) of Table 1.3 show the estimated effects. The estimated labour income response is substantially larger for mothers with more dependent children, with an increase of £258.0 a month (s.e. 107.6) compared with £150.5 (s.e. £60.42) for mothers with fewer children. The extensive margin is similarly larger for mothers with more dependent children, at 5.7 percent (s.e. 6.2) compared with 3.5 percent (s.e. 2.1) for those from smaller families. This is consistent with parents responding to increased salience of future benefit reductions, in addition to the contemporaneous reduction, at the time their first child becomes ineligible. Also consistent with this, the estimated labour supply effects at the time subsequent children age out of eligibility are numerically small and statistically insignificant.

Panel B expresses these treatment effects as elasticities. The myopic elasticities for large families are very large, with an income elasticity of -2.46 and an extensive margin elasticity of -0.55 (although this latter is not statistically significant). The elasticities for smaller families are numerically smaller, at -0.60 and -0.06 respectively, but still large compared to estimates from existing literature (e.g. McClelland and Mok, 2012). The forward-looking elasticites, which assume that the labour supply responses take account of all future CTC reductions as well as the contemporaneous one, are larger for larger families than for smaller families. And, as in the previous subsection, they are smaller than the myopic

[19] The median mother in group L has two dependent children. Around 27 percent have only one and 29 percent have three or more.

elasticities and more in line with existing literature.

Finally, Panel C expresses the estimates as income replacement rates. The myopic replacement rates at the time a family's first child ages out of CTC eligibility are very high for both large and small families. Mothers with large families increase labour income by 163 percent of the contemporaneous reduction in CTC, while those with smaller families replace almost 100 percent. However, the forward-looking replacement rates are smaller at 35 percent for large families and 45 percent for small families. The replacement rates estimated for other children are numerically small and statistically insignificant, reflecting the estimated treatment effect for this group.

As before, the very high myopic replacement rates suggest that the labour supply adjustment is in response to more than just the contemporaneous loss in benefit income. This points to the increased salience of age-related eligibility rules, leading mothers to realise that benefit income will fall further, driving the labour supply response. The similarity of the forward-looking replacement rates between large and small families further supports this conclusion, indicating that the difference in labour supply response between these groups is indeed driven by differences in future benefit reductions.

5.4 Alternative mechanisms

The results set out above are consistent with mothers learning of the rules governing eligibility to CTC through experience. I now discuss a number of alternative mechanisms which could also generate a labour supply response to the foreseeable reduction in CTC, and argue that none is consistent with all features of my results. I therefore conclude that learning about non-salient eligibility rules through experience is the most likely mechanism to have generated my results.

Time inconsistency. Even if eligibility rules are fully salient, there may still be income effects at the time of a foreseeable lump sum benefit reduction if workers are time inconsistent. Time inconsistent workers would consume more than a worker with exponential discounting in the period before losing eligibility for CTC, and so fail to accumulate sufficient savings to smooth marginal utility when they become ineligible. They may therefore reduce consumption and increase labour supply at the time they become ineligible, even if the reduction in their benefits is lump-sum. A substantial literature since Laibson (1997) has documented time inconsistent decision-making in the context of saving decisions.

However, a model of time inconsistency would not predict the asymmetry between first and subsequent children documented in Section 5.2. Instead, time inconsistency would predict a labour supply response at the time all children become ineligible, not just for the first.

Myopia. Myopic workers fail to take account of any future changes in their benefit entitlement.[20] Such workers are distinct from those who learn through experience because they never foresee benefit reductions associated with their children becoming ineligible for CTC, even if they have experienced them before. If workers were myopic, the income effect on their labour supply induced by their child ageing out of CTC would be invariant to whether they have children who became ineligible previously, and would also not depend on the number of children who will become ineligible in future. Both of these predictions are rejected by the results.

Misattribution. Following Liebman and Zeckhauser (2004), a growing literature has documented evidence that, after experiencing a change to their tax or benefits, people may misattribute the cause of the change. For example, a literature has documented cases where people seem to incorrectly believe that a lump sum change in benefits or charges is actually due to an increased withdrawal rate or price (e.g. Feldman et al., 2016; Ito, 2012). However, if parents misattributed the lump-sum change in CTC to a change in their marginal tax rate, but were otherwise well informed about the timing and magnitude of the reduction, there would be no contemporaneous income effect and they should reduce (or, at least, not increase) labour supply when they become ineligible due to a perceived substitution effect. There would be no offsetting income effect if the household anticipated the magnitude and timing of the reduction. The labour supply increases I have documented are therefore inconsistent with such an explanation.

Credit constraints. Even in a standard lifecycle model, labour supply may respond to a foreseeable lump-sum change in unearned income if workers have only limited access to credit, reducing their ability to borrow to smooth consumption. If a worker expects a future lump-sum *increase* in unearned income, their optimal plan would be to smooth marginal utility by borrowing against the future windfall; if they are prevented from doing so, there may be a discontinuous increase in consumption and reduction in labour supply at the time of the windfall.

[20] As such, when defined this way, myopia is a form of extreme time inconsistency in which all future periods are discounted arbitrarily heavily.

Such an argument does not apply to a setting where people anticipate a future income *reduction* (rather than an increase), such as the CTC. In this case, a forward-looking worker does not need to borrow in order to carry out their optimal plan, but instead would accumulate savings in the period before the anticipated lump-sum loss and use these to smooth marginal utility at the time of the loss. Any constraints on borrowing do not bind when attempting to implement such a plan, and so would not be capable of generating the income effects I document in Section 5.

6 Robustness

6.1 Exogeneity

An important assumption in the analysis above is that a child's eligibility for CTC is exogenous to their parents' labour supply. This is plausible as a child's eligibility is primarily determined by their age. However, for children aged between 16 and 20, whether their family remains eligible for CTC depends on whether they stay in full-time education. This introduces potential endogeneity: if the child's decision about whether to continue in education is related to unobservable determinants of their mother's labour supply, A_{it} will be correlated with the residual in equation (1.8) and the resulting estimated labour supply effects will be biased. For example, if a child is more likely to leave education (and so lose eligibility for CTC) if their mother experiences a negative labour supply shock, the estimated responses would understate the true labour supply effect of a reduction in CTC.

While there is widespread agreement that parental income is a strong determinant of children's post-compulsory education choices,[21] evidence on the effect of shocks to parental income is more mixed. However, Carneiro and Heckman (2003) present evidence from the US that short-term family income constraints affect the education choices of only a small fraction of high school graduates, with longer-term family effects as the primary drivers of education choice. This point is corroborated in UK data by Chevalier and Lanot (2002), who find only a limited effect of current family income on the educational attainment of their children. These findings suggest that the education choices of children in my sample are unlikely to be endogenous to transitory parental income shocks.

To provide direct evidence that my findings are not driven by endogeneity in the timing of CTC ineligibilty, I repeat each analysis using the child's age as an instrument for their eligibility. This strategy is designed to use only the exogenous

[21]See e.g. Blanden and Gregg (2004) for a review of evidence from the US and UK.

variation in eligibility arising from child ageing for identification, and to discard the effects of any individual-specific education choices which may be related to innovations in maternal labour supply ϵ_{it}. I implement this instrumental variables strategy using two stage least squares with a first stage linear probability model for CTC ineligibility A_{it} of the form

$$A_{it} = \alpha + \beta_1 \mathbb{1}\left[\text{age}_{it} > 16\right] + \beta_2 \mathbb{1}\left[\text{age}_{it} > 20\right] + \beta_3 (\mathbb{1}\left[\text{age}_{it} \in (16, 20)\right] \times \text{age}_{it}) + X'_{it}\zeta + \theta_t + u_{it} \tag{1.9}$$

and use the predicted values $\hat{A}_{it}^{IV}$ in place of A_{it} in each of the three analyses described in Section 5. I present the estimates for each analysis in Panel A of Table 1.4.[22] While these IV estimates of the labour supply effects are less precise than the baseline estimates in Section 5, reflected in the higher standard errors around the estimates in this section, the magnitude and pattern of the results are very similar.

The estimates in column (1) show that, pooling all children together, losing eligibility for CTC leads mothers to increase labour earnings by £114.6 a month, and to increase the probability of working by 1.6 percentage points. These estimates are similar in size to, and not statistically different from, the £117.3 a month and 2.1 percentage point baseline estimates in column (1) of Table 1.3.

Next, columns (2) and (3) in Panel A of Table 1.4 show that the IV labour supply estimates are concentrated on the first child, with an increase in labour earnings of £181.4 a month and an extensive margin response of 4.0 percentage points for mothers whose first child ages out of eligibility. Again these estimates are close to the baseline estimates of the increase in labour income of £158.9 a month and 3.7 percentage points for first children from Table 1.3. The IV estimates for subsequent children in column (3) are negative but close to zero – and similar to the statistically insignificant baseline estimated responses for subsequent children in Table 1.3.

Finally, the IV estimates in columns (4) to (6) show that the labour supply effect is substantially larger for mothers with more children yet to age out of eligibility. Mothers from larger families increase their labour income by £396.7 a month, and increase the probability of working of 10.8 percentage points, compared to the baseline estimates of £335.4 and 6.7 percentage points in Table 1.3. The IV estimated labour supply responses for mothers with fewer remaining

[22] The sample for this analysis includes only children who turn 20 during the period covered by the sample, rather than those who become ineligible for CTC during the sample period as in Section 5. All children become ineligible for CTC at age 20 regardless of their education choices and so this is the age that generates entirely exogenous variation in CTC.

Table 1.4: Robustness Tests

	All	First Child	Other Child	First Child, Large Fam.	First Child, Small Fam.	Other Child
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Instrument for treatment timing A_{it}						
Labour Income	114.6	181.4	-24.75	239.7	179.8	-17.40
(£/mo.)	(157.2)	(157.3)	(167.7)	(283.9)	(158.6)	(165.8)
Probability of	0.0166	0.0404	-0.0332	0.142	0.0339	-0.0326
working	(0.0516)	(0.0534)	(0.0627)	(0.144)	(0.0526)	(0.0579)
Observations	62,754	62,754	62,754	62,754	62,754	62,754
Panel B: Instrument for treatment group T_i						
Labour Income	125.2	151.7	71.82	338.3	140.4	121.3
(£/mo.)	(119.1)	(130.8)	(156.5)	(256.7)	(127.3)	(149.1)
Probability of	0.0837*	0.119**	0.0377	0.164	0.0820	0.0806
working	(0.0455)	(0.0572)	(0.0751)	(0.106)	(0.0518)	(0.0654)
Observations	72,531	72,531	72,531	72,531	72,531	72,531
Panel C: Differences in probability child leaves home						
Prob. child	0.00294	0.00290	0.00290	0.00551	0.00494	0.00101
leaves home	(0.0142)	(0.0149)	(0.0149)	(0.0290)	(0.0149)	(0.0147)
Observations	72,531	72,531	72,531	72,531	72,531	72,531
Panel D: Differences in pre-period trends						
Labour Income	1.204	0.154	-1.799	1.899	-0.749	2.916
(£/mo.)	(3.113)	(3.358)	(3.714)	(6.757)	(3.585)	(3.497)
Probability of	0.00234**	0.00110	0.00270*	0.00184	0.00134	0.00320***
working	(0.000985)	(0.00115)	(0.00139)	(0.00274)	(0.00116)	(0.00111)
Observations	43,209	43,209	43,209	43,209	43,209	43,209
Panel E: Including single mothers						
Labour Income	163.4***	190.6***	95.4	288.47***	184.90***	96.81
(£/ month)	(59.60)	(61.37)	(63.94)	(93.30)	(61.74)	(61.16)
Probability of	0.033*	0.043**	0.007	0.042	0.044**	0.006
working	(0.018)	(0.019)	(0.025)	(0.057)	(0.019)	(0.023)
Observations	109,042	109,042	109,042	109,042	109,042	109,042

Notes: Table contains estimates for all robustness tests described in Section 6. Standard errors, clustered at level of the mother, in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

children in column (5) are smaller at £166.2 and 3.6 percentage points, corresponding to the smaller baseline estimates of £147.5 and 3.6 percentage points in Table 1.3. And, as before, the results for subsequent children from any sized family are substantially smaller in both the IV and baseline analysis.

These results provide evidence that none of the three key results presented in Section 5 is driven by endogeneity in the timing of CTC ineligibility. While the IV estimates are marginally larger than the main estimates, the differences are small and not statistically or economically significant. And, crucially, the pattern of results remains unchanged: there is a substantial income effect associated with losing eligibility for CTC; the income effect is driven by responses among mothers who lose eligibility for the first time, with a substantially smaller response for subsequent children; and the effect is largest for mothers with more remaining dependent children. The results remain consistent with the rules governing CTC being initially non-salient for claimant families, who then learn about them through experience.

6.2 Treatment group assignment

As I describe in Section 4.2, I divide the sample into control and treatment groups according to their location on the CTC schedule, based on combined household income in the 24 months before their child becomes ineligible for CTC. There are two potential concerns with this approach. First, any measurement error in household income may lead individuals to be assigned to the wrong group. This could attenuate the estimated labour supply response to a lump sum reduction in CTC if individuals who should belong to the control group are assigned to the treatment group, or it could exaggerate the response if the lump-sum treatment group also contains individuals who experience a reduction in their withdrawal rate when their child becomes ineligible.[23]

A second concern relates to transitory income shocks. If a household experiences an income shock in the 24 months before losing eligibility for CTC, this may change the group they are assigned while also generating a predictable future change in labour supply due to mean reversion. For example, a family which would ordinarily be in the control group may enter one of the treatment groups due to a negative labour supply shock in the months before their child becomes ineligible for CTC – and the difference-in-differences procedure would

[23] As I describe in Section 4.2, the treatment group assignment procedure I use is designed to prevent this second form of group mis-assignment, potentially by increasing the chances of the first. The approach is therefore conservative as it works against my main finding of a positive labour supply response.

attribute the subsequent mean reverting increase in their labour supply to losing eligibility for CTC. This could lead the estimates to overstate the true labour supply response to losing CTC.

To provide evidence on whether either of these concerns affects the results, I instrument for treatment group assignment using a form of grouping estimator motivated by Blundell et al. (1998). The key idea is that mothers who have similar exogenous observable characteristics are also likely to experience similar changes to CTC when their children become ineligible. Therefore, even if an individual mother is assigned to the incorrect treatment group due to idosyncratic measurement error or labour supply shocks, the treatment group assignment of similar mothers provides information about the group she should be assigned to.

I assign mothers to mutually exclusive groups based on their exogenous characteristics: age in five categories (29 or younger, 30-39, 40-49, 50-65, and older than 66), highest qualification (less than compulsory education, compulsory education only, post-16 education, first degree, higher degree), and geographical location (based on which of 12 government office regions they live in). I then estimate first stage linear probability models for each variable in $W_{it} = \{T_i^g, T_i^g \times A_{it}\}$ of the following form

$$W_{it} = \alpha + G_i'\phi_1 + (G_i \times A_{it})'\phi_2 + X_{it}'\zeta + v_{it} \tag{1.10}$$

where G_i is a vector of dummy variables indicating which of the mutually exclusive groups the mother of child i belongs to. I then use the predicted values as instruments in each of the three main estimating equations.

I present the estimated labour supply effects for each of the three main analyses in Panel B of Table 1.4. As with the IV analysis for eligibility status, the magnitude and pattern of the point estimates are very similar to the baseline estimates across all three analyses, albeit less precise as reflected in higher standard errors. The results in column (1) show that the large overall labour supply response to losing eligibility for CTC is robust to instrumenting for treatment group assignment. Columns (2) and (3) show that the estimated labour supply responses remain substantially larger for mothers who lose eligibility for the first time, compared with those who have already had a child age out of eligibility, supporting the hypothesis of learning through experience. And the results in columns (4) to (6) show that the finding of larger labour supply effects for mothers with more remaining children than those with fewer is also robust to this IV strategy.

The estimates in this section show that, as with the IV results discussed in the previous subsection, the three key results in Section 5 are not driven by endogeneity in treatment group assignment. The pattern of results in Panel B of Table 1.4 is the same as the main results in Section 5, and remain consistent with CTC claimants learning about initially non-salient eligibility rules through experience.

6.3 Changes in household composition

A separate concern is that the results may be driven by differential changes in household composition at the time the child becomes ineligible for CTC. For example, if there are differences between the control and treatment groups in the likelihood a child leaves home when they become ineligible for CTC, this would invalidate the identification assumption that the direct effect of child ageing on labour supply is the same between these groups.

I test whether this is the case by repeating each of my three main analyses with a binary variable for whether the child leaves home as the dependent variable. The estimates of $\{\beta_g\}$ from this analysis, which I report in Panel C of Table 1.4, indicate whether children in treatment group g are more likely to leave home than those in the control group when they become ineligible for CTC. These estimates are close to zero and statistically insignificant across all analyses, indicating that the results are not driven by differential changes in household structure.

6.4 Parallel trends

An important assumption in the difference-in-differences empirical design is that the direct effect of a child becoming ineligible for CTC on their mother's labour supply is the same across the treatment and control groups. This is simply a parallel trends assumption, with the time dimension defined as child age relative to the eligibility threshold for CTC, which ensures that the control group provides an unbiased estimate of counterfactual labour supply for those in the treatment groups.

While it is not possible to test this parallel trends assumption directly, I examine whether it holds in the period before becoming ineligible for CTC. Figure 1.2 provides initial graphical evidence that this is the case. This shows that the labour supply of those in the lump-sum treatment group L evolves similarly to those in the control group C in the months before losing eligibility for CTC, but diverges sharply in the period after.

I now test formally for different linear trends in labour supply between the control group and each of the treatment groups. I estimate regression equations of the following form on the 24 month period before the child loses eligibility:

$$y_{it} = \alpha + \eta t + \sum_g \beta_g (T_i^g \times t) + \sum_g \gamma_g T_i^g + \zeta X_{it} + \theta_t + \epsilon_{it}. \tag{1.11}$$

where η is the linear trend in monthly labour supply y_{it} for control group members in the 24 months before losing eligibility for CTC, and β_g is the difference in linear trend between those in treatment group g and those in the control group. Panel D of Table 1.4 presents the estimates of β_g for both labour income and the binary labour supply variable. All estimates for the labour income variable are economically small and statistically insignificant, indicating that the trend of labour income is not different between the control and any of the treatment groups in the 24 months before losing eligibility for CTC.

The estimates for the binary labour supply variable are also economically small, but do indicate some statistically significant differences. However, the estimates in columns (2) and (3) suggest that there is no statistical difference in pre-trends for mothers who experience a reduction in eligibility for the fist time; instead, all of the difference in pre-trends is concentrated on mothers whose subsequent children become ineligible. This result is corroborated by the estimates in columns (4) to (6).

I find this reassuring for two reasons. First, the labour supply effects I document in Section 5 are concentrated among mothers who lose CTC for the first time – and there is no evidence of differential pre-trends for this group. Indeed, if the trend in the probability of work is higher for subsequent children than the control group, this would work against my finding of a null effect on labour supply. Second, any higher trend in the treatment group for subsequent children than in the control group could simply reflect responses to these mothers' first child becoming ineligible, and is therefore consistent with the model of learning through experience.

Overall, I find no evidence that my main results are driven by differential trends in labour supply between the control and treatment groups.

6.5 Including single mothers

In the main analysis I focus on the labour supply of mothers who are married or cohabit with a partner. The reason for this restriction was to use spousal labour income as a source of variation in the effects of a child ageing out of eligibility for CTC: mothers with similar incomes, but whose partner earn different amounts, would experience a different change to the CTC schedule. However, single mothers are an important focus of welfare policy, and it is interesting to consider whether the salience effects I document for married mothers are also present for this group.

In Panel E of Table 1.4 I show each of the main estimates for the entire sample, including both married and single mothers. The pattern of results is unchanged relative to the analysis for married mothers in Table 1.3, with point estimates which are marginally higher but very similar in magnitude. The salience of the eligibility rules therefore does not appear to be substantially different for single mothers.

7 Conclusion

This chapter provides new evidence that a dynamic feature of a major welfare programme are non-salient, leading claimants to make substantial optimisation errors due to their failure to anticipate future benefit reductions. The results also suggest that claimants learn about the eligibility rules through experience, and hence better anticipating benefit reductions after their first experience with them.

The results underline the potential for incentives to be non-salient, even in real-world systems with high financial stakes. Finding that claimants learn about the eligibility rules through experience suggests that the effects are most likely to be present in systems people interact with infrequently. This leaves open the possibility that salience effects are present in other complex environments with infrequent interaction. This motivates my study of pension incentives in Chapter 3.

Appendix to Chapter 1

1 Data Appendix

In this appendix, I provide further details on the way I constructed the data used in Chapter 1. I use data drawn from the UK Household Longitudinal Study (UKHLS) linked to it predecessor the British Household Panel Survey (BHPS). BHPS was a representative panel survey of UK households conducted annually from 1991 until the end of 2008; UKHLS has followed a larger sample of households, including remaining BHPS sample members alongside new recruits, since 2009. I restrict attention to 2003 to 2013, corresponding to the period from the introduction of CTC to the start of its phased replacement.

The surveys contain detailed information on (i) demographic characteristics for all household members including age, highest educational qualification, and region of residence (ii) labour supply for adult respondents, including gross labour earnings and employment status; (iii) income from all other sources, including benefit payments and (iv) education status for children, including whether the child is still in education and, if not, when they left. The data provide links between family members, so it is possible to identify a mother's children and spouse or cohabiting partner (if she has one).

1.1 Construction of monthly data

While each sample member is interviewed only at annual increments, respondents are asked to report any employment changes they have experienced since their previous interview. The result is detailed information about each sample member's employment status at the time of their interview once a year, alongside recalled information about earnings and employment status for any jobs which may have been missed between interviews. I these data to create a panel of earnings and employment status at monthly frequency for all sample participants. This allows me to focus on the months directly surrounding families losing eligibility for CTC in estimating any labour supply response.

There are two steps to constructing the monthly data from the annual survey. First, I create a list of all spells of employment (or non-employment) reported by sample members, including any spells reported to have started and ended in the period between interviews. For each job, I record whether the individual was employed part-time or full-time, their gross labour earnings, and the months that the job started and ended (if the individual is not still employed in the job at the time of their most recent interview).[24]

Second, I merge the employment spells onto a monthly panel using the start and end dates. As has also been noted by for BHPS and UKHLS by Smith (2011), this is complicated by a number of inconsistencies in the dates individuals claim to have started or ended employment spells. There are broadly two categories of inconsistency:

Inconsistencies with start and end dates. In some cases, the date an individual reports ending one spell of employment, unemployment or inactivity does not match the date they report starting the next. This results in either a gap in the individual's employment history or a period where the spells overlap. I have used the following procedure to assign start and end dates to avoid these inconsistencies:

a. Set the start date of each individual's first employment (or non-employment) spell equal to the date of their first interview. I therefore ignore any employment recalled from the period before the individual joined the sample.

b. Set the start date of all subsequent employment (or non-employment) spells equal to the reported end date of the previous spell.

c. Set the end date of the an individual's most recent employment spell equal to the date of their most recent interview.

Inconsistencies with information provided in different waves. In some cases, individuals provide contradictory information about a particular part of their job history in different waves. For example, they may report different transition dates between employment spells, or different earnings information, or different answers to whether they were part-time or full-time. In these cases, I have given preference to information provided in interviews conducted most

[24]These data are available for all employment spells recorded in the BHPS data, and for employment spells reflecting the respondents' current status at the time of their interviews in UKHLS. However, UKHLS did not collect earnings data for spells which started and ended between interviews. I drop spells with missing earnings data in the analysis, but note that these missing data affect only 59 mothers.

recently to the employment spells being asked about – i.e., older interviews. This is equivalent to the "closest interview method" discussed by Smith (2011).

1.2 Definition of key variables

In this section I describe the ways I have constructed key variables used in the analysis.

Labour supply. I define monthly labour income as the gross usual pay an individual receives from all jobs they hold, including any profits from self-employment. I set monthly labour income to zero for any months an individual reports to be unemployed or economically inactive. For the binary labour supply variable, I determine an individual to be working if (1) they report being either employed or self-employed in a given month and (2) they report non-zero labour income.

Tax credit eligibility. I have calculated each family's entitlement to tax credits – i.e. both CTC and WTC – based on their reported characteristics in each month. I calculate entitlement to WTC as well as CTC because the two tax credits are paid together and assessed jointly against household income. A family's entitlement to WTC affects the level of total tax credits it is entitled to, but does not affect the change in payments induced by child ageing dependent status.

The eligibility calculations depend on:

- the combined household income of a mother and her co-resident partner Y;
- the number of dependent children K, where a child is classed as dependent if they are (i) co-resident with their parents, (ii) aged younger than 16, or younger than 20 and in full-time education; and
- the number of hours each parent works each week. This determines whether the family is entitled to WTC and, if so, how much they can claim.[25] The amount of CTC a family is entitled to does not depend on the hours worked by parents.

I describe the way CTC entitlement is determined in Section 2. WTC entitlement is determined by family structure, subject to adult members working a minimum number of hours a week. People in couples (which all in my sample are) can claim WTC if

[25] I treat a part-time job as 20 hours a week, and full-time as 40 hours.

1. they have dependent children and work at least 24 hours a week between them, with one member of the couple working at least 16 hours; or

2. if they do not have dependent children and either member of the couple works at least 30 hours a week.

Eligible couples received an average of £3,496 a year over the period 2003-2013, with an additional £719 available to those who jointly worked at least 30 hours. WTC is means tested against combined household income, and is withdrawn at rate τ_1 above some threshold $\bar{Y}_0$, which is strictly lower than the threshold above which CTC starts to be withdrawn $\bar{Y}_1$.[26]

I note that a family moves from category (1) to category (2) when their final child is no longer classed as dependent. While this could provide an incentive for parents to increase their hours if neither works at least 30 hours a week, the actual impact of child ageing out of dependent status on WTC entitlement is negligible. Estimating regression equation (1.8) with WTC entitlement as the dependent variable reveals that the actual impact of child ageing out of dependent status on the WTC entitlement of group L is just $-£0.51$ a month (s.e. £0.38). This is both statistically and economically insignificant, and so I ignore changes to WTC in my analysis: children ageing out of dependent status only affects a family's CTC claim.

Treatment group assignment. In Section 4.2 I describe the three main groups used in the analysis: L, M and C. I assigned mothers to treatment groups using the following procedure:[27]

1. For each of a mother's children, in each of the 24 months t before the child ages out of eligibility, calculate:

   (a) the *actual* level of CTC she is entitled to and the withdrawal rate she faces, $b(K_t, Y_t)$ and $\tau(K_t, Y_t)$;

   (b) the *counterfactual* level of CTC she would have been entitled to and the withdrawal rate she would have faced if she had one fewer dependent child, $b(K_t - 1, Y_t)$ and $\tau(K_t - 1, Y_t)$;

   (c) the *impact* of her child ageing out of CTC on her level of benefits and

[26] Over the period 2003 to 2013, $\bar{Y}_0$ increased from £5,060 to £6,420.

[27] A mother could potentially be assigned to different treatment groups for each of her children if the impact of the children becoming ineligible on her CTC schedule is different.

withdrawal rate as the difference between the two, i.e.

$$\Delta b(K_t, Y_t) := b(K_t, Y_t) - b(K_t - 1, Y_t)$$
$$\Delta \tau(K_t, Y_t) := \tau(K_t, Y_t) - \tau(K_t - 1, Y_t).$$

This method of calculating the impact of CTC hold household income Y constant at the level it was prior to losing eligibility, and isolates the impact of a reduction in the number of dependent children on the CTC schedule.

2. Assign mothers to treatment groups for each of their children based on the *impact* of that child ageing out of eligibility on the CTC schedule. I used the following procedure:

   (a) assign a mother to the control group C if $\Delta b(K_t, Y_t) = \Delta \tau(K_t, Y_t) = 0$ for each of the 24 months t before losing eligibility;

   (b) assign a mother to the lump-sum treatment group L if $\Delta b(K_t, Y_t) \neq 0$ in any of the 24 months t before losing eligibility, but $\Delta \tau(K_t, Y_t) = 0$ for each of the 24 months t before losing eligibility;

   (c) assign a mother to the marginal tax rate treatment group M if $\Delta \tau(K_t, Y_t) \neq 0$ in any of the 24 months t before losing eligibility.

   Given the interest in the chapter in mothers who experience a purely lump-sum reduction in CTC, this procedure is conservative in assigning mothers to group L. I assign a mother to group M if there is even one month in the two years before losing eligibility in which the mother's withdrawal rate would have changed had her child aged out of eligibility in that month. This is to ensure, as far as possible, that the labour supply response I estimate for group L is truly in response to a lump-sum benefit reduction.

Chapter 2

Benefit Salience: A Structural Model

1 Introduction

In Chapter 1, I provide empirical evidence the the rules linking entitlement to the UK's Child Tax Credit to the age of a family's children are imperfectly salient, but that families learn about the rules through experience. However, there are a number of important and interesting questions which cannot be answered directly from the empirical results due to the fundamentally dynamic nature of the setting. Specifically, how widespread is the non-salience of the eligibility rules? Are the average labour supply responses driven by relatively few claimants making a substantial labour supply adjustment, or from smaller adjustment from a higher proportion of claimants? And what are the welfare costs of non-salient eligibility rules in this dynamic setting?

To answer these questions, I develop a structural life-cycle labour supply model in which households may be unaware that part of their benefit income depends on the age of their children. I estimate key parameters of the model by indirect inference, matching the empirical results to identify the proportion of claimants who are uninformed of the rules separately from preference parameters determining their labour supply responsiveness. The estimated preference parameters are closely in line with existing literature, underlining that my empirical results are consistent quantitatively with introducing non-salient eligibility rules to an otherwise standard labour supply model. I also estimate that 82 percent of claimants are unaware of the benefit rules, suggesting widespread non-salience of the policy feature.

Finally, I assess the consequences for welfare. I calculate the reduction in initial assets households would be willing to accept, on average, to be aware of the eligibility rules in all periods. This compensating variation in initial assets is substantial—equal to 14 percent of the lifetime present value of CTC. And I show that, in this setting, there are no offsetting benefits to the government from non-salient rules. This points to major source of inefficiency in the welfare system which has not previously been documented. Overall, the results underline the importance of considering the salience of dynamic policy features, as they may be highly non-salient (even when the financial stakes are high), generating substantial welfare costs.

The contribution of this chapter is to incorporate non-salient dynamic incentives into an otherwise standard structural life-cycle labour supply model. I use the model both to validate my empirical results in Chapter 1, showing that non-salient eligibility rules are capable of generating the responses quantitatively as well as qualitatively. I then use the model to assess the welfare implications of the non-salient rules. The results highlight that the lack of salience of the eligibility rules leads to life cycle labour supply and savings decisions which are substantially different from those people would make if they were fully informed. This underlines that the effects I document in Chapter 1 are important to consider when designing policy.

The rest of this chapter is organised as follows. I specify the model outlined in Chapter 1 fully in section 2. In section 3, I explain how features of the empirical results from Chapter 1 provide sufficient variation to identify key parameters of the model and present my estimates. I use the estimated model to assess the welfare costs of the non-salient eligibility rules in section 4, before summarising and offering conclusions in section 5.

2 Model

In this section, I specify the model outlined in Chapter 1 fully, with parameterised wage and employment equations and a specific utility function. This allows me use the empirical results from 1 to estimate key parameters within the model and assess the welfare implications of non-salient eligibility rules.

2.1 Overview of the model

Households enter the model at age 20, and choose how much female labour to supply and how much to consume (rather than save) each year until retiring with certainty at age 60. After retirement, the household faces exogenous mortality risk and consumes from accumulated savings for a maximum of 20 further years.

Each household contains a cohabiting adult couple and up to five children. I assume that adult household members retire and die at the same age.

Throughout its working life, the household faces three sources of risk: (i) innovations to the female wage, (ii) innovations to the male wage and (iii) whether the male household member is employed. I model explicitly the main features of the UK's progressive tax and transfer system which provides partial insurance against these sources of risk. Households may also experience a shock to expected lifetime wealth when their children age out of eligibility for CTC if the rules governing their eligibility were not salient.

I make a number of further simplifications to reduce the computational burden of solving the model. First, I assume that the age gap between children is deterministic (equal to two years) for all families. This means I only need to keep track of the age of one child (the oldest) rather than the age of all children separately. Second, because the motivation for the model is labour supply responses among married women with older children, I do not consider marriage, divorce or endogenous fertility.

2.2 Model details

In each year between entering the labour market at age 20 and retiring at age 60, a household chooses how much female labour to supply (n) and how much to consume (c) to maximise expected lifetime utility, taking as given its current circumstances. These states are its age (t), accumulated assets (a), female wage (w), male wage (w^m), male employment status (n^m), the total number of its children ($\bar{k}$), the age of its oldest child (t^k) and its awareness of rules linking benefit entitlement to the number and age of its children (θ). I represent the set of state variables at age t as X_t.

Preferences

Utility is separable over time, but is non-separable between consumption and female labour supply within periods. Instantaneous utility is given by

$$u(c_t, n_t) = \frac{\tilde{c}_t^{1-\gamma}}{1-\gamma} \exp\{U(n_t)\}, \tag{2.1}$$

where $\tilde{c}$ is equivalised family consumption and n is female labour supply which can take one of three values: not working (O), part-time (P) or full-time (F).[1] The parameter γ determines both risk aversion and the elasticity of intertemporal

[1] The consumption equivalence scale s is $1.6 + 0.1k$, and equivalised consumption is related to total household consumption c by $\tilde{c} = c/s$.

substitution and the function U, which determines how utility from consumption is affected by working, is given by

$$U(n_t) = \mathbb{1}\,[n_t = P]\alpha_P + \mathbb{1}\,[n_t = F]\alpha_F. \tag{2.2}$$

Since I estimate (1-γ) to be negative in Section 3 below, positive U for full- or part-time work means that working reduces utility. This specification also means that labour supply and consumption are complements. The household's problem at age t can be written as

$$V_t(X_t) = \max_{\{\tilde{c}_j, n_j\}_{j=t,\ldots,\bar{t}}} E_t \Big\{ \sum_{j=t}^{\bar{t}} \beta^{j-t} u(\tilde{c}_j, n_j) | X_t \Big\} \tag{2.3}$$

subject to the budget constraint,

where the expectation is taken over future random variation in female wages w, male wages w^M and male employment n^M given the current state X_t. In particular, this expectation is conditional on the current understanding of how benefit income will evolve in the future as captured by $\theta_t \in X_t$.

Budget constraint

I define the household's budget constraint in terms of the asset evolution equation

$$a_{t+1} = (1+r)a_t + w_t n_t + w_t^m n_t^m + T(n_t, X_t) - c_t, \tag{2.4}$$

$$a_{t+1} \geq \underline{a}_{t+1}, \tag{2.5}$$

with the initial condition $a_0 = 0$ and terminal condition $a_{\bar{t}+1} = 0$. Households can save or to borrow up to the natural borrowing constraint, i.e. the maximum borrowing that the household can repay with certainty, regardless of how stochastic processes realise in future.[2] I describe each component of the asset accumulation equation below.

Female wages. Women earn income $w_t n_t$ from working n_t hours at wage w_t, where n_t can take the values 0, 72 or 152 hours a month. The female wage process is given by

[2] I have also considered a version of the model in which households are credit constrained and are unable to borrow. This has only minor effects on the model estimates.

$$\ln w_t = \alpha_0 + \alpha_1 \ln(t-20) + v_t, \tag{2.6}$$

$$v_t = \rho v_{t-1} + \epsilon_t, \tag{2.7}$$

where v_t is an AR(1) individual productivity process with iid innovations $\epsilon_t \sim N(0, \sigma^2)$. Wages follow an age profile, reflecting the effects of accumulated labour market experience in a computationally parsimonious way.

Male earnings. The woman's partner earns income $w_t^m n_t^m$, where both w_t^m and n_t^m are exogenous. Men either work 160 hours a month (F) or not at all (O), with their wage and employment status following exogenous processes given by

$$\ln w_t^m = \alpha_0^m + \alpha_1^m \ln(t-20) + v_t^m, \tag{2.8}$$

$$v_t^m = \rho^m v_{t-1}^m + \epsilon_t^m, \tag{2.9}$$

$$\text{Prob}(n_t^m = F) = \text{Prob}(\omega_t^m > m_1 + m_2 \mathbb{1}\left[n_{t-1}^m = F\right]) \tag{2.10}$$

where v_t^m is an AR(1) individual productivity process with iid innovations $\epsilon_t^m \sim N(0, \sigma_m^2)$. Male employment also follows an autoregressive process with innovations $\omega_t^m \sim N(0,1)$. I follow Blundell et al. (2016a), who find no evidence of selection in male employment in the UK, in specifying independent processes for male earnings and employment.

Net transfers. The household receives transfers (net of taxes) $T(n_t, X_t)$. I model the tax and transfer system, as well as the household's entitlement to major welfare programs, from the 2010/11 tax year in the UK.[3] This year was typical of the period from 2003 to 2013, which is the period I study in the empirical analysis.

The component of the net transfer function I focus on is the household's entitlement to CTC. As described in Section 2, a household's CTC entitlement in period t depends the number k_t of its $\bar{k}$ children who are classed as *dependent*. The number of dependent children is given by

$$k_t = \bar{k} - g(\bar{k}, t, t^k). \tag{2.11}$$

[3]In particular I model the household's income tax and National Insurance liabilities, along with potential entitlement to Child Tax Credit, Working Tax Credit, Jobseekers' Allowance, and Child Benefit. Over the period I study, Child Benefit was substantially less generous than Child Tax Credit and, crucially, was not means tested.

where $g(\bar{k}, t, t^k)$ are the eligibility rules. These determine how many of a family's $\bar{k}$ children are ineligible for CTC in year t based on the age of the oldest child, t^k. The total amount of CTC a household can claim in period t is $b(k_t, Y_t)$, which depends on joint household income and the number of dependent children as given by equation (1.2) in Section 2.

Salience of CTC eligibility rules

The household may not be aware of how the number and age of its children determine eligibility for CTC. The *salience* of these child-related eligibility rules $g(\bar{k}, t, t^k)$ is given by $\theta \in \{0, 1\}$. If $\theta = 1$ the household knows how its entitlement depends on its children's characteristics, but if $\theta = 0$ it incorrectly believes that its entitlement is unrelated. A household with θ thinks that its number of dependent children in period t will be

$$\tilde{k}_t(\theta) = \bar{k} - \theta g(\bar{k}, t, t^k) \tag{2.12}$$

Non-salient eligibility rules will therefore lead a household to over-estimate the level of benefit income it will receive in periods where its actual benefit eligibility will be reduced.

3 Structural Analysis

The results in Chapter 1 imply that there are CTC recipients who fail to anticipate the benefit reductions which will arise from their children becoming ineligible, but subsequently learn through experience with the system. However, a number of important questions are difficult to answer directly from the empirical results. Specifically, what proportion of claimants are initially unaware of the rules linking eligibility to the age of the their children? How many learn through experience? And what are the welfare implications?

In this section, I use the empirical results from Chapter 1 to estimate these key parameters within the structural life-cycle labour supply model outlined in Section 2. I then use the model to compute the welfare cost of being uninformed about the benefit eligibility rules over the life-cycle, along with the implications for the government budget.

3.1 Estimation

Target moments and identification. Within the structure of the model set out in Section 2, the empirical results identify preference parameters (which determine the size of the labour supply response to an unanticipated reduction in benefits) separately from the proportion of claimants who fail to anticipate the benefit reduction. These features have important implications for the welfare analysis.

I estimate three preference parameters $\{\gamma, \alpha_P, \alpha_F\}$ and the average salience of the eligibility rules $\bar{\theta}_0$ by indirect inference, matching four empirical features: (i) average earnings in the period before becoming ineligible, (ii) the female employment rate in the period before becoming ineligible, (iii) the relative labour income response of large vs small families, and (iv) the average labour supply response to the first child becoming ineligible for CTC.

The identification argument is simple and intuitive. While all moments jointly provide identifying variation for all parameters, (i) and (ii) provide information on the costs of full- and part-time work α_P and α_F; (iii) provides variation in the response to differently sized future income reductions necessary to identify the elasticity of intertemporal substitution $1/\gamma$; and, given knowledge of the preference parameters, (iv) provides information on the average number of households who were uninformed about the benefit reduction $(1 - \bar{\theta})$ as a scaling factor between the model-generated labour supply response (under the assumption that *all* households are unaware of the benefit eligibility rules) and the empirical responses.

Parameters set exogenously. I use a two-step procedure to estimate the model, first setting a number of parameters exogenously based on existing literature or characteristics of my sample. These are the annual discount and interest rates, parameters of the male and female wage equations, and male employment transitions. I list the values of these parameters in Table 2.1.

The wage equation parameters draw on Blundell et al. (2016a), who estimate male and female wage processes using the same dataset, and for a similar period, as my empirical work. I show the life-cycle profiles of average wages in Figure 2.1.

I estimate the parameters of the male employment process for partners of women in the lump-sum treatment group – which is the focus of both my empirical work and the model – using a probit regression of a binary employment variable on its (annual) lag. The employment process implies an average male employment rate of 74.9 percent, compared to 74.2 percent in the sample.

I also set family composition outside the model, based again on the lump-sum treatment group in my sample. When I simulate the model, I allow for families to have between one and five children in total, and for the first child to age out of eligibility when the household is age 40, 45 or 50. I show composition of the

Table 2.1: Parameters Set Outside the Model

Parameter	Symbol	Value
Economic Environment		
Annual discount factor	β	0.98
Annual interest rate	r	0.015
Female wages		
Initial mean log wage	γ_0	1.85
Coefficient on log age	γ_1	0.08
St. dev. of innovation	σ	0.22
Autocorrelation coefficient	ρ	0.91
Male wages		
Initial mean log wage	γ_0^m	1.94
Coefficient on log age	γ_1^m	0.09
St. dev. of innovation	σ^m	0.12
Autocorrelation coefficient	ρ^m	0.97
Male employment		
Constant	m_0	-1.25
Coefficient on lag employment	m_1	3.00

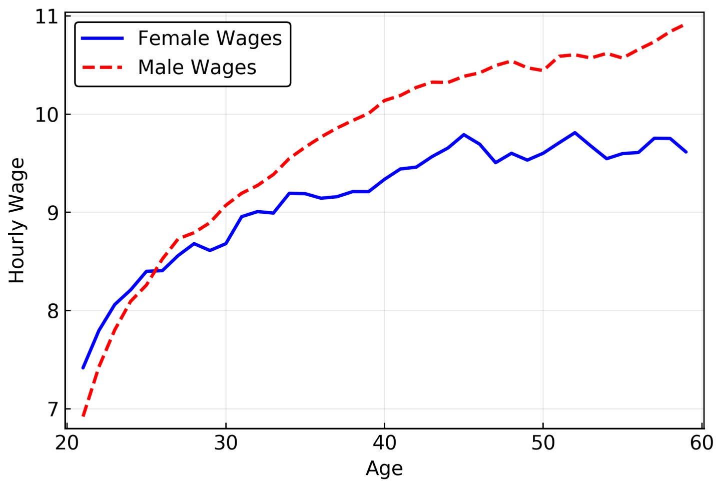

Figure 2.1: Male and Female Wage Processes

Table 2.2: Simulation Sample Composition

	Age when first child ages out			
Number of children	40	45	50	**Total**
1	1.7%	2.7%	10.6%	**15.0%**
2	13.4%	15.7%	17.9%	**47.0%**
3	9.9%	7.9%	5.2%	**23.0%**
4	4.5%	3.2%	2.3%	**10.0%**
5	2.3%	1.9%	0.9%	**5.0%**
Total	**31.8%**	**31.4%**	**36.7%**	**100.0%**

simulated sample in Table 2.2.

These proportions are set so that the simulated sample matches the lump-sum treatment group L in the sample, and captures the empirical relationship between family size and age of children. Specifically, the proportions of households by number of children are equal to the proportions in the sample. Then, within each family size, I calculated the proportion of households in which the mother was aged less than 42.5, between 42.5 and 47.5, and older than 47.5 when their first child became ineligible. I assign these groups to 40, 45 and 50 in the model.

Estimation procedure. I then estimate the four remaining parameters by indirect inference. These are three preference parameters $\{\gamma, \alpha_P, \alpha_F\}$ and the average salience of the eligibility rules $\bar{\theta}$. First, I simulate life-cycle profiles of consumption and labour supply for a large number of households under the assumption that they are perfectly informed about the benefit eligibility rules. I then simulate the same households again, this time under the assumption that they are unaware of the benefit eligibility rules but become informed after their first child becomes ineligible. I hold the values of all random processes constant between the two sets of simulations.

This procedure allows me to calculate the labour supply response to learning about the benefit eligibility rules through experience for each household, using the same household's behaviour in the full salience case as the counterfactual – an ideal experiment within the model. For each household i, I calculate the change in female labour income at the time child k becomes ineligible under the full salience and learning through experience scenarios: respectively, $\Delta y_{i,k}^{S}$ and $\Delta y_{i,k}^{NS}$. Household i's labour supply response to the benefit reduction when child k becomes ineligible is given by

$$\Delta y_{i,k} = \Delta y_{i,k}^{NS} - \Delta y_{i,k}^{S}. \tag{2.13}$$

I then average these household-specific treatment effects. I construct treatment groups analogously to the empirical section, and compute the average labour supply response within each of these groups generated by the model. I assign each household i to group L, M or C using an identical procedure to that described in Section 4.2, and further split group L into first and subsequent children, and large and small families. The average labour supply response in group $g \in \{L_{first}, L_{first,large}, L_{first,other}\}$ is given by

$$\Delta\bar{y}_g = \frac{1}{N_g} \sum_{\{i,k\}\in g} \Delta y_{i,k} \tag{2.14}$$

where N_g is the number of household-child pairs $\{i, k\}$ in group g.[4] This is the average labour supply response if *all* households in group g fail to anticipate the benefit reduction. The average labour supply response if only proportion $(1-\bar{\theta})$ of households fail to anticipate the reduction in CTC is $(1-\bar{\theta})\Delta\bar{y}_g$, as the response among the fraction $\bar{\theta}$ of households which anticipate the reduction is zero. I match $(1-\bar{\theta})\Delta\bar{y}_g$ for $g \in \{L_{first}, L_{first,large}, L_{first,other}\}$ to the empirical counterparts reported in Section 5.

I also match the average employment rate and earnings for group L in the two years before their first child becomes ineligible. Note that average earnings and employment are higher for households who anticipate the reduction in CTC in the period before the first child becomes ineligible. I construct average earnings in the model as the weighted average of earnings from households in salience and non-salience scenarios in the two years before the first child becomes ineligible,

$$\bar{y}_g = (1-\bar{\theta})\bar{y}_g^{NS} + \bar{\theta}\bar{y}_g^{S}, \tag{2.15}$$

for $g = L_{first}$, and equivalently for the average employment rate $\bar{e}_g$.

3.2 Parameter estimates and model fit

Table 2.3 shows the values of the estimated preference parameters and proportion of claimants initially unaware of the benefit eligibility rules. I estimate that 82 percent of CTC claimants are initially unaware of the rules governing eligibility. A substantial proportion of claimants therefore fail to anticipate the reduction in benefit income when their first child ages out of ineligibility. To the best of my knowledge, this is the first estimate of the proportion of claimants unaware of a dynamic policy feature with such substantial financial stakes.

[4] As in practice I only match moments at the time a family's first child becomes ineligible, N_g is simply equal to the number of households in group g.

Table 2.3: Estimated Parameters

Parameter	Symbol	Value
Proportion initially unaware	1-$\bar{\theta}$	0.82
Coeff. Relative Risk Aversion	γ	2.24
Fixed cost full-time work	α_F	0.61
Fixed cost part-time work	α_P	0.31

It is unclear, ex ante, whether the average level of salience in my setting would be higher or lower than found by existing literature in other contexts (typically sales taxes) (e.g. Abeler and Jäger, 2015; Chetty et al., 2009; Feldman et al., 2018; Taubinsky and Rees-Jones, 2017). While the financial stakes are substantially higher in the setting I study – which may provide an incentive for claimants to become well informed – the welfare system in general is much more complex, which may stymie learning. It is notable that the proportion of uninformed claimants I estimate is high, and indicates the importance of considering potentially non-salient incentives in real-world welfare systems, even when the financial stakes are high.

Table 2.4: Model Fit

Moment	Target Value	Model Value
Pre-period employment rate	0.578	0.533
Pre-period average earnings	526.1	514.6
Large family / small family response	1.714	1.916
Average response at first child	158.9	151.4

The estimated preference parameters are in line with existing literature. The coefficient of relative risk aversion γ is within the range of values estimated by Blundell et al. (1994), Attanasio and Weber (1995) and subsequent literature. And the costs of full- and part-time work are consistent with the estimates from Blundell et al. (2016a) for low-educated married mothers. The consistency of these preference parameters with existing literature underlines that adding non-salient eligibility rules with learning to an otherwise standard life-cycle model is capable of generating the empirical results quantitatively as well as qualitatively.

I compare the targeted moments to those generated by the model in Table 2.4. Generally the model does well in matching the target moments. Again, this indicates that the mechanism of initial non-salience of benefit eligibility rules with subsequent learning is capable of generating quantitatively similar results to the empirical analysis.

Finally, in Figure 2.2 I show the lifecycle profiles of female labour supply and household consumption, under the full salience and learning scenarios. The rows of the figure differ by the age of the household when the eldest child becomes ineligible: the household first loses eligibility for CTC at age 40 in panels (a) and (b), at age 45 in (c) and (d) and at age 50 in (e) and (f). I note three main features. First, labour supply is lower in the learning scenario than in the full salience scenario over the period before the first child becomes ineligible, but this pattern reverses in the period after. This labour supply increase in the learning scenario compared with full salience reflects the information shock at the time the first child becomes ineligible. And the reverse is true for consumption, with a reduction in the learning scenario compared with full salience at the time the first child becomes ineligible.

Second, there are predictable increases in labour supply and reductions in consumption even in the full salience scenario. This solely reflects changes to the consumption equivalence scale at the time the child ages out of eligibility, and captures any direct effects of child ageing on the labour supply of their parents. This highlights the importance of having a credible control group in the empirical section, designed to hold these predictable labour supply increases constant.

Finally, the responses to labour supply and consumption are larger for households who are younger when the first child becomes ineligible. This is for two reasons. First, as shown in Table 2.2, these younger households have a higher composition of larger families. These families have a larger information shock in the learning scenario and so inrcease labour supply by more. Second, households who lose eligibility for CTC when they are relatively young experience a larger information shock, as there are more years in which the actual entitlement to CTC is below their initial (mistaken) expectation.

4 Welfare implications

I now consider the welfare implications of the non-salient eligibility rules, using compensating variation as the welfare metric. I solve for the reduction in initial assets ΔA_0 which reduces expected lifetime utility for those in group L in the full-salience scenario (EV_0^S) to the level experienced under non-salience and learning (EV_0^{NS}), allowing lifecycle consumption, savings and labour supply choices to adjust:

$$EV_0^S(\Delta A_0) = EV_0^{NS}(0) \tag{2.16}$$

I show the compensating variation in row (1) of Table 2.5, separately by family size and as a weighted average across all families. Larger families experience

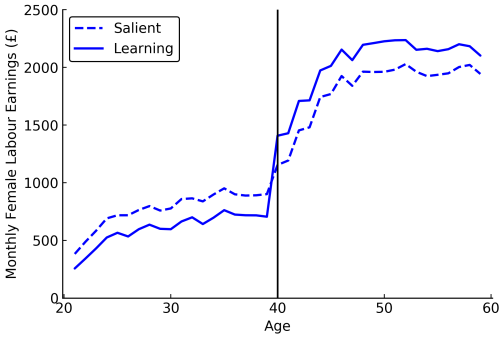

(a) Female Labour Income (lose CTC age 40)

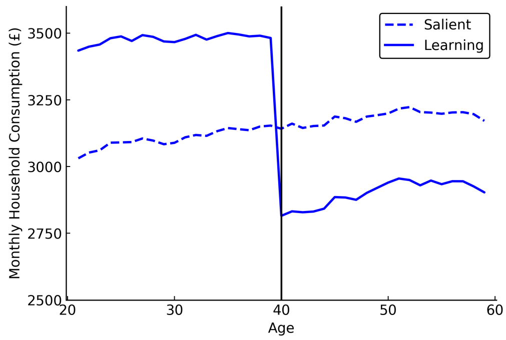

(b) Household Consumption (lose CTC age 40)

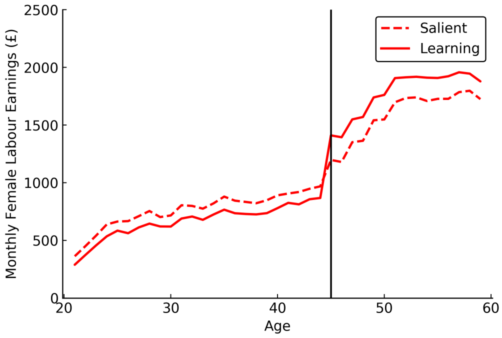

(c) Female Labour Income (lose CTC age 45)

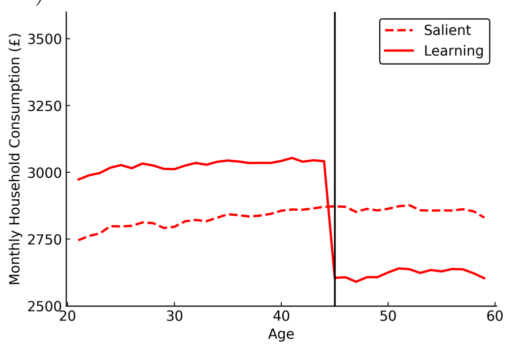

(d) Household Consumption (lose CTC age 45)

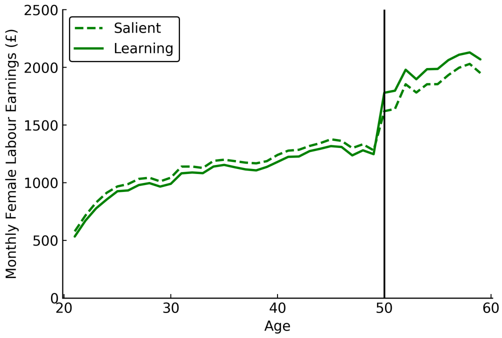

(e) Female Labour Income (lose CTC age 50)

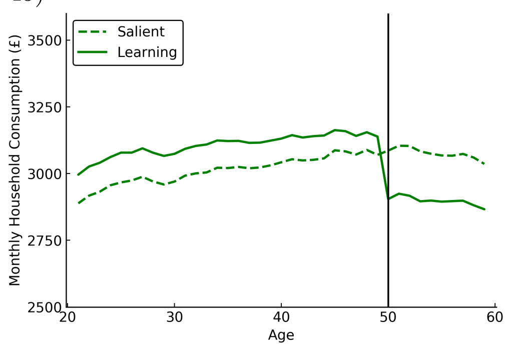

(f) Household Consumption (lose CTC age 50)

Figure 2.2: Lifecycle Profiles of Monthly Labour Supply and Consumption

higher welfare costs from being unaware of the eligibility rules than smaller families because the size of the information shock is higher. On average, a household would be willing to pay £18,412 at age 20 to be perfectly informed of the benefit eligibility rules rather than uninformed. In row (2), I scale the compensating variation by $(1 - \bar{\theta}) = 0.82$, the proportion of claimants I estimate to be uninformed about the eligibility rules. I interpret this as the is the welfare cost incurred by an average household.[5]

To provide a measure of the size of the welfare cost, I express average welfare cost as a proportion of total lifetime CTC receipt. In row (3), I show the present value of lifetime CTC receipts for families of different sizes, and in row (4) I express the average welfare costs as a proportion of these lifetime CTC receipts. On average, a family would be willing to accept 13.8 percent cut in their CTC payments in exchange for being informed about the eligibility rules. This underlines that the welfare costs I estimate are substantial.

In some situations, a government may benefit from obfuscating financial incentives. If so, the benefits to government may offset, in part or in full, the welfare costs borne by individuals – and could in principle make non-salient features part of optimal policy design. Commonly, the benefits to government stem from reducing the behavioural response to a distortionary tax by reducing its salience (e.g. Chetty et al., 2009). However, such a motivation is not present in the situation I study, as the existence of a foreseeable and lump-sum reduction in benefit income is not distortionary.

Nonetheless, it could be that the government would benefit from claimants being uninformed if this yields higher tax revenue which could be used to reduce other distortionary taxes. As I show in Figure 2.2, the life-cycle profile of labour supply depends strongly on households' awareness in the eligibility rules. Given the progressive tax and transfer system in the UK, the total net revenue collected by the government from households who are unaware of the eligibility rules (whose increase in labour supply could push into a higher tax bracket may be higher than from those who are informed of the rules) may be higher than those who are unaware. Whether or not the net effect on government revenue is positive or negative is unclear ex ante.

There is no benefit to government revenue in this setting. I show the present value of tax revenue (net of transfer payments) collected by government over the lifecycle for households who are fully aware of the eligibility rules in row (5), and for those who are initially uninformed in row (6). One-child households are net contributors to the government budget over their life-cycle, but households with two or more children are net recipients. This reflects that the group I study are particularly low income families and, by construction, are eligible for child-related benefits.

In column (7) I calculate the difference in net government revenues between

[5] In fact, this is likely a lower bound on the welfare cost, as it assumes that informed households incur no cost from the potentially non-salient eligibility rules. However, it may be that these households bear other costs associated with becoming informed about the benefit eligibility rules which are not reflected in my model.

Table 2.5: Welfare

		Family Size					
		Avg.	1	2	3	4	5
Compensating Variation							
(1)	Compensating ΔA_0	-18,412	-7,977	-14,582	-23,423	-30,873	-38,272
(2)	$(1-\bar{\theta}_0)\Delta A_0$	-15,098	-6,541	-11,957	-19,207	-25,316	-31,383
Expressed as a proportion of time-zero present value of CTC							
(3)	PV of CTC receipts	109,480	42,193	85,423	135,431	193,555	253,020
(4)	(2)/(3)	-13.8%	-15.5%	-14.0%	-14.2%	-13.1%	-12.4%
Effect on Government Budget							
(5)	Net receipts, full salience	-90,702	10,508	-56,157	-128,552	-214,330	-300,200
(6)	Net receipts, learning	-92,914	8,626	-58,505	-131,119	-216,376	-300,887
(7)	(6) − (5)	-2,212	-1,881	-2,348	-2,567	-2,047	-687

the full salience and learning scenarios. The government's budget deteriorates by an average of £2,212 per household, in present value terms, if the household is unaware of the benefit eligibility rules. This arises because average labour supply is lower for households who are unaware of the eligibility rules in their early working life, leading to lower tax revenue and higher welfare payments. These added costs to the government are not offset (in present value terms) from the higher net revenue later in life. This underlines that, in this case, effects on the government budget cannot justify obfuscating the eligibility rules.

The results therefore point to a major source of inefficiency in the welfare system which has not previously been documented, but which is important to consider when designing and assessing welfare policy.

5 Conclusion

In this chapter, I develop and estimate a structural model of life-cycle labour supply in which the rules linking benefit entitlement to the age of a family's children may not be salient. There are two key findings. First, using the empirical results from Chapter 1 as targets, I identify the proportion of claimants who are unaware of the eligibility rules separately from other preference parameters. The estimated preference parameters are in line with existing literature. Conducting an ideal experiment within the model—comparing the behaviour of the same households across cases where they are and are not aware of the eligibility rules—closely replicates the empirical results. This suggests that my proposed mechanism of initial non-salience with subsequent learning is capable of generating my empirical findings quantitatively as well as qualitatively.

Second, I provide what is, to the best of my knowledge, the first estimate of the average salience of a dynamic policy feature which generates such substantial variation in income. I find that the proportion of households initially unaware of the CTC eligibility rules is high, despite the size of the financial stakes. Finally, I compute the welfare cost to claimants from being uninformed about the eligibility rules. The costs are substantial—equal to nearly 14 percent of the lifetime present value of CTC, with no offsetting benefits to government revenue.

The results underline the importance of considering the salience of dynamic policy features when designing and assessing welfare policy. Even when the financial stakes are high, the evidence in this chapter shows that the features may be highly non-salient, generating substantial welfare costs. This points to major source of inefficiency in the welfare system which has not previously been documented.

Chapter 3

Pension Salience

1 Introduction

Saving for retirement is one of the most important decisions people make over their working life. However, the rules governing the value of retirement saving schemes, and the consequent incentives they provide for saving and labour supply over the life cycle, are often complex. If individuals do not fully understand these incentives, or restrict attention to a subset which are particularly salient, their resulting decisions about how much to work and save may be suboptimal. While there is considerable policy interest in the extent to which individuals are financially prepared for retirement, direct evidence on whether people are aware of features of their pension portfolio—and how this varies across pensions with different characteristics—is more limited.

In this chapter, I study the salience of incentives to continue working (rather than retire) provided by private pension schemes in the UK. There are three reasons these incentives may be imperfectly salient. First, pensions are complex. Their value may be a function of the individual's earnings and employment history, contributions made each year (either themselves or by their employer), annuity prices, and potentially stock market returns over their working lives. This complexity is compounded in settings, such as the UK, where it is common for individuals to hold multiple pensions, as the rules governing the value of each may be quite different.

Second, and related, pension incentives are heterogeneous between schemes. This is particularly true in the UK system, which has high coverage of private pensions (which differ in their features) in addition to a relatively small state pension. This makes it difficult for individuals to learn about features of their own pension arrangements from peers or generic information. And, for individuals holding multiple pensions, knowledge of the features of one scheme is often insufficient to fully understand their entire pension portfolio.

Finally, the feedback individuals receive from their pension is limited, and often arrives late in working life. This precludes learning about pension features through experience.[1] And, at the time many individuals do find out about their pension entitlement (either at retirement or just before),[2] there is often limited scope to remedy any shortfalls in retirement saving. This pension shock is a problem for policy to solve—and developing a better understanding of its causes is important for designing effective policy responses.

In this chapter, I provide empirical evidence consistent with differential salience of the financial incentives to continue working provided by private pensions. As in the first two chapters of this thesis, my approach is to study labour supply decisions to shed light on the salience of the underlying incentives which influence them. I exploit the substantial heterogeneity in private pension schemes in the UK to estimate the labour supply response to similarly-sized incentives arising from different features of the pension system.[3] If people are equally well informed of all features of their pension portfolio, the labour supply responses should be similar. But if some incentives to continue working are less salient than others, the labour supply responses to the non-salient features would be muted. Differential salience of the incentives would therefore translate into differently sized labour supply responses to similarly-sized financial incentives.

My data are drawn from the first seven waves of the English Longitudinal Study of Ageing (ELSA), a representative panel survey of individuals aged 50 and over living in England between 2002 and 2017. A particular benefit of the ELSA data is that it contains detailed information on pension arrangements including—for each pension an individual is contributing to, holds retained rights for, or is drawing—the value of the pension and any features which will govern changes in its value over time or with continued work. I use this information to construct forward-looking measures of the incentive to continue working provided by each individual's specific pension arrangements.

I have three main sets of results. First, I estimate the labour supply response

[1]The evidence I present in Chapter 1 of this thesis suggests that learning through experience is an important feature of the way claimants of the UK's Child Tax Credit come to understand its dynamic features.

[2]For example, Rohwedder and Kleinjans (2006) document that individuals' knowledge of the value of their pension increases as they approach retirement.

[3]In this chapter I distinguish between incentives provided by DB and DC pensions, but in principle my approach could be applied to more granular definitions of pension "feature" – such as distinguishing between employee and employer contributions to a DC pension plan, or distinguishing increments to a DB pension due to scheme tenure or earnings.

to the combined financial incentives provided by all pensions an individual holds. I follow a large existing literature by estimating option value-type models of retirement timing (following Stock and Wise, 1990), which relate the decision about whether to continue working to the difference between the utility from retiring immediately and retiring at the optimal (utility-maximising) date. I find a positive and statistically significant impact of marginal incentives to work on continued labour supply. These findings are consistent with, and provide an update to, existing evidence for the UK (e.g. Banks et al., 2016; Blundell et al., 2002).

A criticism of the option value framework is that retirement decisions are posited as a function of both pension incentives and wages—and differences between individuals in wages generate a substantial fraction of the identifying variation.[4] This introduces endogeneity to the estimates if both wages and retirement timing are driven by unobserved heterogeneity in tastes for work. I deal with this in two ways: (i) I also use an alternative measure of financial incentive which does not include wages directly, based on Coile and Gruber's (2007) "peak value" model,[5] and (ii) I exploit the panel structure of my data to restrict the identifying variation to within-individual changes in pension incentives over time. I show that labour supply responses to pension incentives remain positive and significant (albeit smaller than in the traditional option value model). This reassures that estimates of labour supply responses to pension incentives in existing literature are not driven solely by endogeneity arising from unobserved tastes for work.

Second, I test for differential labour supply responses to similarly-sized retirement incentives arising from DB and DC pensions. To motivate this approach I first show that, while the labour supply incentives provided by DB pensions tend to be sharper on average than for DC pensions, there is still substantial overlap in the size of the incentives provided by each type of pension—with variation for each type across individuals and as a function of age. I then estimate separate labour supply responses to incentives provided by DB and DC pensions and show that labour supply responds more strongly to incentives provided by a DB pension. In my preferred specification (which includes an individual fixed effect, restricting the identifying variation to within-person changes over time), a £1000 increase in the peak value of a DB pension increases the probability the individual is in work between ages 50 and 70 by 0.077 percentage points.

[4] I show that this is the case in the ELSA data in section 3.

[5] As I discuss in section 3, an added benefit of the peak value model for my application is that the measure is linear in pension wealth. It is therefore possible to estimate responses separately between incentives provided by each of an individual's pensions in a way the non-linear option value does not permit.

This is statistically significantly higher than the 0.011 percentage points for a DC pension, and both responses are statistically significantly different from zero.

The stronger response to incentives for DB pensions is consistent with the schemes typically being less complex: the value of a DB pension is generally a function only of a measure of late-career earnings, tenure with the scheme, and age relative to a threshold specified in the scheme. By contrast, the value of a DC pension depends on the individual's and employer's contributions, the value of the fund, and investment returns and annuity prices at retirement. And so, despite being capable of generating incentives to continue work similar in size to a DB pension, the higher complexity of DC plans may prevent individuals from fully understanding their features. This is a potentially important consideration for policymakers, especially given recent trends towards DC pension plans (Blundell et al., 2016b).

Finally, I test for differences in the size of labour supply responses for individuals holding different numbers of private pensions—a measure of the complexity of their pension portfolio. There is considerable heterogeneity in the number of private pensions people hold in England: the median member of my analysis sample holds two private pensions, but over 10 percent of the sample hold five or more. I find that the labour supply responses are attenuated for individuals holding more pensions—even when using only within-person variation in incentives for identification. This suggestive evidence is consistent with the salience of labour supply incentives provided by pensions decreasing with the complexity of an individual's portfolio.

Related literature. This chapter is at the intersection of two distinct literatures: first on the determinants of retirement decisions, and second on people's potential unawareness of financial incentives.

My overall approach is to study the effects of the financial incentives provided by pension schemes on labour supply, exploiting the heterogeneity in pension across people (and over time) in the UK. A large literature has documented that financial incentives provided by pension arrangements determine individuals' decisions about whether (or how much) to work at older ages (e.g. Banks et al., 2016; Blundell et al., 2002; Coile and Gruber, 2007; French, 2005, among others). I provide updated estimates on the responsiveness of individuals in England to the incentives provided by private pensions. And, in particular, I adopt an approach designed to mitigate potential endogeneity arising in option value retirement models from both wages and retirement timing being influenced by unobserved tastes for work. I show that labour supply responses to private pension incentives remain positive and statistically significant, reassuring that existing findings in the option value literature are not driven solely by endogeneity.

This chapter also contributes to literature documenting that individuals are imperfectly informed about aspects of their pension portfolio. For example, Gustman and Steinmeier (2004) document that people have limited knowledge of their pension and social security wealth. And Lusardi and Mitchell (2007) show that differences in financial literacy account for variance in the degree of retirement planning.

A number of papers document that providing workers with information about their pension arrangements can improve their knowledge of their incentives and affect their labour supply. Mastrobuoni (2011) shows that sending letters with information about social security benefits has a significant impact on workers' knowledge about their future benefits in the US, despite the information being available previously. However, he finds no effect of this improved knowledge on retirement timing. Liebman and Luttmer (2015) send older workers information about their social security benefits in a randomised field experiment, and show that the intervention increases labour supply. Dolls et al. (2018) study the effect of similar information letters in the German pension system, and show that they led to an increase in both pension savings and labour supply.

Overall, this existing literature provides evidence that individuals are imperfectly informed about features of their pension—and, particularly, the incentives they provide to continue working. This chapter contributes by documenting differential labour supply responses across features of an individual's pension portfolio, consistent with heterogeneity in the salience of incentives across different pension schemes. And, by testing for differences in understanding using variation between pensions (rather than studying the effects of a particular policy, such as the information interventions studied in existing literature), my empirical strategy attempts to identify differential salience from revealed behaviour in the absence of a policy intervention. It therefore avoids potential concerns of confounding behavioural responses to a specific intervention with effects of increased information.

A number of papers document excess sensitivity to certain highly-salient pension features—notably, age thresholds. Seibold (2019) documents an excess mass of exits from employment at statutory retirement ages in Germany, despite there being little financial incentive to do so. Kim (2020) studies retirement decisions around age-related milestones which generate differential increases in pension wealth across individuals in the Missouri public school system. He shows that retirement timing is strongly influenced by the milestones, but not by the extent to which pension wealth is increased. Similarly, Cribb et al. (2016) document excess sensitivity to the state pension age in the UK, exploiting gradual increases to the female state pension age. MacCuish (2019) studies the same reform, arguing that the effects are driven by a model of rational inattention.

These findings on excess retirement at age milestones are consistent with individuals being imperfectly informed of rules determining their pension benefits, but the age milestones being particularly salient features of the system: people focus on the salient age milestones, but not on the less salient financial incentive they provide.

More generally, this chapter contributes to the literature on the salience of financial incentives more generally (Abeler and Jäger, 2015; Chetty et al., 2009; Taubinsky and Rees-Jones, 2017). As in the first two chapters of this thesis, the setting I study in this chapter is distinct from much existing literature in that the environment I study (i) is intrinsically dynamic, as individuals do not experience the consequences of behaviour over the life-cycle on pension value until retirement, and (ii) has very high financial stakes. Failing to consider all aspects of their pension portfolio could therefore lead individuals to make large optimisation errors, generating considerable welfare cost.

The rest of this chapter is organised as follows. I start by providing an overview of the pension system in the UK in section 2. I then explain my empirical strategy for identifying labour supply responses in section 3 and provide an overview of the ELSA data, along with how I construct key variables, in section 4. I then present my main results in section 5. I provide concluding remarks, and suggest directions for future research, in section 6.

2 Institutional Background: Pensions in the UK

There is a substantial variation in pension arrangements across individuals in the UK, owing to the mix of private and state-provided pensions. I provide a brief overview of the main features of the UK pension system here. For more complete descriptions, see Blundell et al. (2002) or Banks et al. (2016).

Most individuals are eligible to receive a relatively small public pension, known as the state pension, when they reach state pension age (historically set at 65 for men and 60 for women, although this has been gradually increasing for women since 2010). The state pension in the UK consists of two parts. The main part, known as the Basic State Pension, depends on the number of years an individual has worked and made National Insurance Contributions.[6] Individuals may also be entitled to the State Second Pension, which is related to career earnings (up to a cap). However, most individuals have historically opted out of this State Second Pension to instead contribute to a private pension in exchange for paying reduced National Insurance Contributions on labour income (Banks et al., 2016). Overall, the state pension provides only weak incentives to leave the labour force. There is no earnings test for the state pension, meaning that entitlement is not reduced if individuals work and claim the pension simultaneously,[7] and those who delay claiming beyond the state pension age are entitled to a generous actuarial uplift.[8]

Owing in part to the historically low level of the state pension, the majority of individuals in the UK are also a member of a private (i.e. non-state) pension scheme. In 2011/12, 60 percent of workers aged 50 to 69 had some form of private pension (Banks et al., 2016). These private schemes typically provide much stronger labour supply incentives than the state pension system, and exhibit considerable heterogeneity in these incentives. They are the focus of this chapter.

Defined Benefit (DB) pensions provide a fixed stream of income in retirement, typically calculated as a function of the individual's late-career earnings and tenure with the pension scheme. Schemes will also typically impose an actuarial reduction for individuals who claim before the scheme-specific normal retirement age, but not provide an actuarial uplift for claiming late. This creates an incentive to claim at exactly the normal retirement age—although the strength of this incentive may vary between schemes.

Individuals with Defined Contribution (DC) pensions, by contrast, pay into an investment fund during working life with the pension value on retirement depending on the total contributions and investment returns. DC pensions therefore typically provide less sharp incentives to retire than DB pensions—although variation in the expected growth of the fund, along with declining annuity rates with age, mean that DC schemes do provide incentives for labour supply.[9]

A number of people in the UK receive additional income from the welfare system at older ages—notably disability insurance. However, Banks et al. (2016)

[6]In practice, owing to allowances made for years an individual is out of work to care for children or disabled adults, or receiving unemployment or disability benefits, the majority of individuals qualify for the full basic state pension by the time they reach state pension age.

[7]Any state pension income is subject to income tax at the individual's marginal tax rate.

[8]Individuals are entitled to a one percent uplift for every five weeks they delay claiming, equivalent to 10.4 percent a year Cribb et al. (2016). However, Crawford and Tetlow (2010) note that only 5 percent of individuals in the UK between state pension age and 75 had deferred receipt of their pension in 2008/09.

[9]I quantify these incentives, and compare them with the incentives provided by DB schemes, in Figure 3.2 of section 4.

note that the majority of the financial incentive to retire is driven by pensions rather than potential eligibility for disability benefits in the UK.[10] I therefore restrict attention to retirement incentives provided by pensions in this chapter.

3 Empirical Strategy

My approach is to estimate labour supply responses to pension incentives by exploiting the heterogeneity in pension arrangements across people (and over time) in the UK. To do this, I construct forward-looking measures of the incentives to work provided by each individual's pension arrangements, by comparing the benefit to an individual from retiring immediately to the benefit they would receive from retiring at some optimal date in the future. My analysis is therefore in the spirit of Stock and Wise (1990) and a large subsequent literature (e.g. Banks et al., 2016; Blundell et al., 2002; Coile and Gruber, 2007) which aims to capture the incentives provided by entire evolution of pension wealth with continued work—as opposed to, say, simply analysing the one-year accrual rate of pension wealth.

The particular focus of my analysis is whether similarly-sized financial incentives to work, produced by different parts of the pension system, lead to differently sized labour supply responses. If individuals are equally well informed about all parts of their pension portfolio, they should respond equally to the same-sized marginal incentive to continue working regardless of which part of the pension system produces it. However, if individuals focus less on certain features of their pension arrangements, any behavioural effect of these less salient incentives may be muted.[11]

I consider two different forward-looking measures of the financial incentive to continue work (rather than retire) provided by an individual's private pension portfolio: the option value as defined by Stock and Wise (1990), and the peak value as defined by Coile and Gruber (2007). I now describe the calculation of each measure, and then discuss their relative suitability for this study.

[10] In particular, they simulate the effects of abolishing disability insurance altogether and conclude that doing so would increase the average number of years worked between ages 50 and 69 by just 0.1.

[11] It could also be that case that individuals *over*-respond to particularly salient features. For example, the widely-documented excess sensitivity to the state pension age in the UK (Cribb et al., 2016), or normal retirement ages in general (e.g. Kim, 2020; Seibold, 2019), may reflect that these features of the pension system are particularly salient. My analysis focuses on identifying differences in labour supply response (reflecting differential salience), but I do not assess whether these differences are due to under- or over-responses. I return to this point in the conclusion.

Option value. Option value models of retirement assume that individuals compare the utility of retiring immediately to the expected utility from retiring at every possible date in the future. For an individual at age t, the expected present value of retiring at age $r \geq t$ is

$$V_t(r) = \sum_{s=t}^{r-1} \beta^{s-t}\pi(s|t)U_W(Y_s) + \sum_{s=r}^{T} \beta^{s-t}\pi(s|t)U_R(B_s(r)) \tag{3.1}$$

where $U_W(Y_s)$ is the value from working with income Y_s, $U_R(B_s(r))$ is the value from being retired with income $B_s(r)$ and $\pi(s|t)$ is the probability of surviving to age s from age t. Following Banks et al. (2016), I assume utility is CRRA and allow the marginal values of income to depend on whether an individual is in work or retired. Under these assumptions, equation (3.1) can be re-written as

$$V_t(r) = \sum_{s=t}^{r-1} \beta^{s-t}\pi(s|t)Y_s^{\gamma} + \sum_{s=r}^{T} \pi(s|t)\beta^{s-t}[kB_s(r)]^{\gamma} \tag{3.2}$$

where I set $k = 1.5$, $\gamma = 0.75$, $\beta = 0.97$. I then compute the option value $OV(t)$ by comparing the value from retiring at each age t to maximum possible value achievable by retiring at any future age

$$OV(t) = V_t(r^*) - V_t(t) \tag{3.3}$$

where $r^* = \arg\max V_t(r)$ is the value-maximising retirement age. This provides a measure of the incentive for an individual to continue working provided by their pensions.

Peak value. Coile and Gruber (2007) construct a measure of "peak value" of pension wealth, which is similar to option value except it (i) excludes labour earnings between time t and the date of retirement r and (ii) does not include utility parameters. A peak value model therefore assumes that individuals compare only the pension wealth they would be entitled to if they retire at different dates when choosing the optimal date to retire.

Specifically, the time-t value of an individual's pension wealth P if they retire in year $r \geq t$ given by

$$P_t(r) = \sum_{s=r}^{T} \pi(s|t)\beta^{s-t}B_s(r) \tag{3.4}$$

and the peak value PV given by difference between retiring today and at pension-maximising age

$$PV(t) = P_t(r^*) - P_t(t) \tag{3.5}$$

Discussion and choice of primary measure. For this analysis, peak value models have two main advantages over option value models. First, as the peak value measure is linear in the present value of pension wealth, it possible to decompose the total financial incentive into the contributions of different pensions (or types of pension): the total peak value an individual faces in a given year is simply the sum of the peak values induced by each of their pensions. Such a simple decomposition is not possible with the non-linear option value measure, making it difficult to estimate separately the effects of different parts of an individual's pension portfolio on their labour supply.

Second, as noted by Coile and Gruber (2007), including labour earnings in the definition of option value means that variation in this measure between individuals is driven largely by variation in earnings.[12] This may introduce endogeneity in the estimation of the impact of financial incentives on continued labour supply if differences in earnings across individuals reflects heterogeneity in unobserved tastes for work. This concern is mitigated for the peak value measure which does not include earnings directly.[13]

For these reasons, my focus is on estimating the impact of financial incentives as measured by peak value. However, to allow comparisons with estimates from option value models in existing literature, I also present estimates of the effect of all private pensions grouped together from option value models.

4 Data Overview

I use the English Longitudinal Study of Ageing (ELSA), a large panel survey representative of the private household population in England aged 50 and over. ELSA contains detailed information about household members' demographic characteristics, labour supply and earnings, and wealth. Crucially for this study, this includes highly detailed information on each household member's pension arrangements. I use data from the first seven waves of ELSA, covering the period from the introduction of the study in 2002 until 2015.[14]

[12]I provide direct evidence that this is the case for my sample in section 3.3 below.

[13]The peak value measure does include the individual's earnings *indirectly* through their effect on accumulated pension wealth. However, as earnings do not enter the definition directly, less variation in peak value is driven by earnings (relative to the particular features of the individual's pensions) than is the case for option value models.

[14]ELSA sample members are interviewed every two years.

The ELSA data contain a total of 72,893 observations on 18,489 individuals over this period. I focus on individuals between the ages of 50 and 70 who are working at the time of at least one of their interviews, and drop 9,810 individuals who do not meet these criteria. I then drop 2,823 individuals because one or more of the covariates required for the analysis is missing, such as accrued state pension or other wealth. My final sample contains 26,316 observations on 6,345 individuals (3,222 men and 3,123 women). In the rest of this section, I describe the key variables required for this study.

4.1 Pension Features and Retirement Incentives

Labour Income

I define labour income Y_t as the sum of earnings from main and any subsidiary employment, net of taxes and National Insurance contributions. For individuals who are still working at the time of their interview in year t, is also necessary to forecast labour income for each year between t and every possible future retirement date r. This is because (i) the value of the pension at date r may depend on the individual's (gross) labour earnings in the years leading up to retirement and (ii) labour earnings between t and r enter the calculation of the option value in equation (3.1) directly. I follow Crawford (2012) and Banks et al. (2016) by assuming that labour income grows in real terms by 2.5 percent a year until age 55, after which earnings remains constant in real terms.

Pension Features

ELSA contains highly detailed information on each household member's pension arrangements. It covers all of an individual's pensions—including those which they are no longer contributing to but are yet to receive an income from, as well as pensions the individual already receives.

For each of an individual's DB pensions, this includes the individual's tenure with the pension scheme and accrual rate (which are required for calculating the value of the pension in the year it is claimed), and indexing rules which govern the value of the pension after it is claimed. Individuals are also asked about any lump-sum they expect to receive from their pension when they retire. For each DC pension, individuals are asked to provide an estimate of the current value of the fund and, for active pension schemes, information on employer and individual contributions.

In addition to this information on private pensions, 80 percent of the ELSA sample consented for their responses to be linked to their history of National

Insurance contributions—which are the main determinant of state pension eligibility. It is therefore possible with ELSA to forecast the value of an individual's state pension and each of their private pensions for each possible future date of retirement.

I merge forecasts of the value of individuals' pensions into the main ELSA dataset.[15] These variables measure, for each possible future year of retirement r, the pension income the individual would receive in each subsequent year s (expressed in constant 2017 prices). They are therefore equivalent to $B_s(t)$ used in the definition of option value (equation 3.1) and peak value (equation 3.4). The variables I use are aggregated across all pensions of the same type—e.g. all DB and all DC pensions.[16]

Survival Probabilities

Each of the three measures of pension incentive described in Section 3 relies on an estimate of the probability the individual survives to age s given that they have survived to age $t \leq s$, $\pi(s|t)$. I calculate these probabilities using actuarial tables for the period 2002-2004, relating to the first wave of ELSA. The tables contain estimates of $\pi(t|t-1)$—the probability of surviving to age t conditional on having reached age $t-1$—separately for men and women. I use these to write the probability of surviving to k as

$$S(k) = \prod_{t=1}^{k} \pi(t|t-1), \tag{3.6}$$

and so the probability of surviving to age s, conditional on having reached age t, is

$$\pi(s|t) = \frac{S(s)}{S(t)}. \tag{3.7}$$

I show the estimated cumulated survival functions $S(k)$ for men and women in Figure 3.1.

Retirement Incentives

I use the data on labour income, pension features and survival probabilities to compute each of the measures of retirement incentive provided by the individual's

[15] See Crawford (2012) for more detail on the construction of these measures.

[16] In principle, the ELSA data would allow for these variables to be constructed at the level of each pension, rather than just each pension *type*. However, I did not have access to the more disclosive pension-level data while preparing this chapter, and so leave study of different responses across individual pensions to future research.

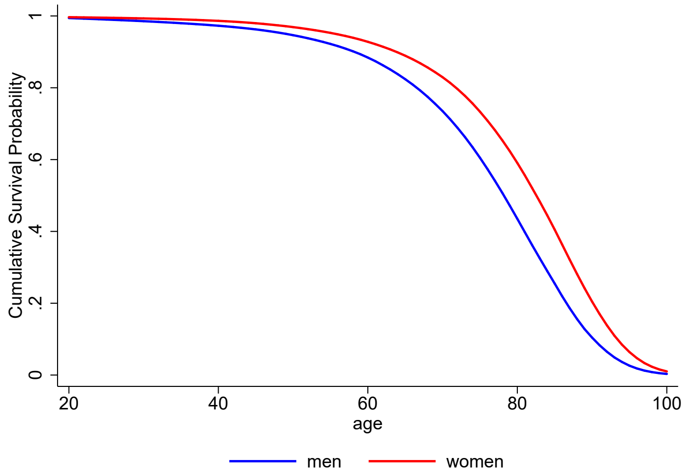

Figure 3.1: Survival Probabilities

pension arrangements described in section 3.

First, in Figure 3.2, I document substantial heterogeneity in the incentives provided by private pensions to continue working (rather than retire). Each panel shows the distribution of $PV(t)$ as defined in equation (3.5), separately by gender and type of pension. These panels show the amount of pension wealth individuals could accumulate by continuing to work until the age that maximises the present value of their pension wealth. For readability, I exclude from the figures individuals whose value of $PV(t)$ is zero (largely reflecting those who do not hold a pension of that type).

While DB pensions provide stronger incentives to continue working on average (as reflected by greater mass further up the distribution of PV), I note that (i) both DB and DC pensions can provide similarly strong incentives to continue working, as measured by potential to accrue pension wealth by continuing to work until the date in future which maximises the present value of pension wealth, and (ii) there is heterogeneity in the strength of these incentives between people. This motivates studying whether individuals respond to similarly-sized incentives to continue working differently depending on whether the incentives are provided by a DB or DC pension.

Second, in Figure 3.3, I show kernel regressions of the retirement incentives provided by DB and DC pensions on age. Panels (a) and (b) use PV as the measure of retirement incentive, while (c) and (d) use option value OV. The difference between DB and DC pensions in the age profile of retirement incentives is clearer for the peak value measures than for the option value measures. This is because the peak value measures exclude income earned between time t and

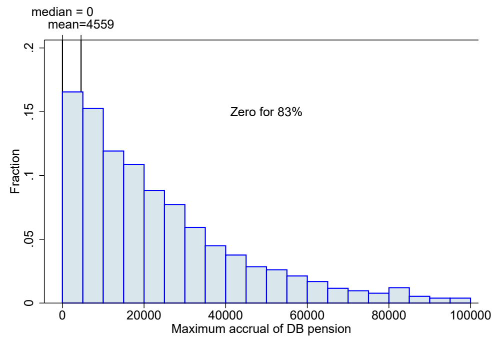

(a) Male maximum DB accrual

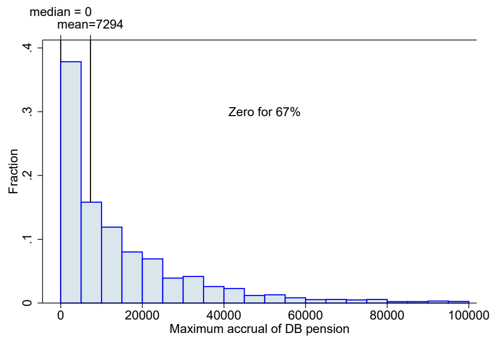

(b) Male maximum DC accrual

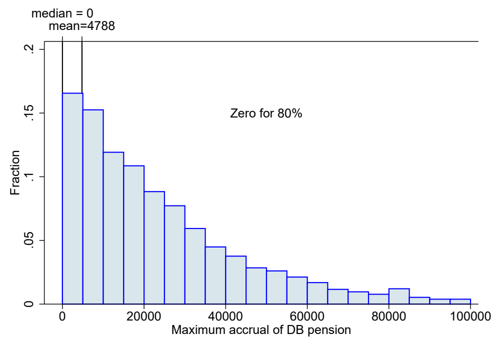

(c) Female maximum DB accrual

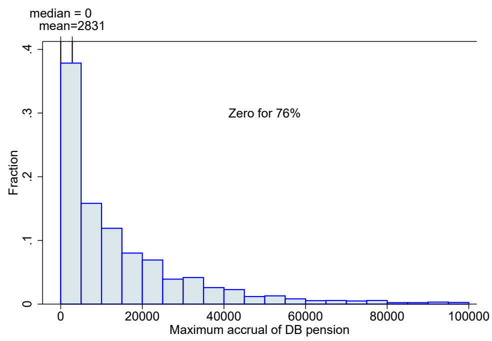

(d) Female maximum DC accrual

Figure 3.2: Maximum accrual of private pension wealth if work until pension-wealth-maximising date

the optimal retirement age and so focus solely on the effect of continued labour supply on pension wealth. By contrast, the option value measure also includes the present value of income earned between t and retirement (which is the same regardless of whether an individual holds a DB or DC pension), obscuring the effect of retirement incentives arising from pension wealth.

As discussed in section 3, the extent of variation in the option value measure arising from earnings introduces potential endogeneity into option value retirement models if individuals with higher unobservable taste for working have both higher earnings and retire later.

4.2 Other variables

I define an individual as being in work in in year t if they define their working status at the time of their interview as "in work" (rather than "not in work"). This is the dependent variable in my analysis.

I also make use of other financial and demographic characteristics available in

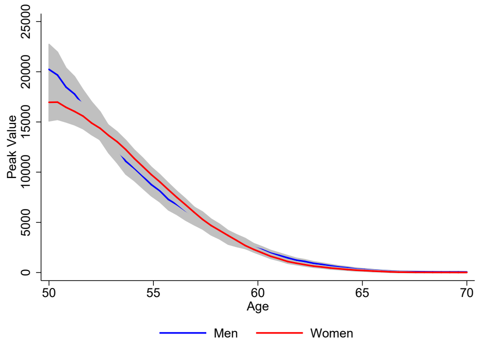

(a) Peak value (DB)

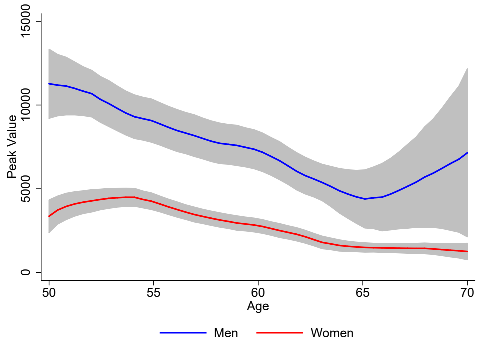

(b) Peak value (DC)

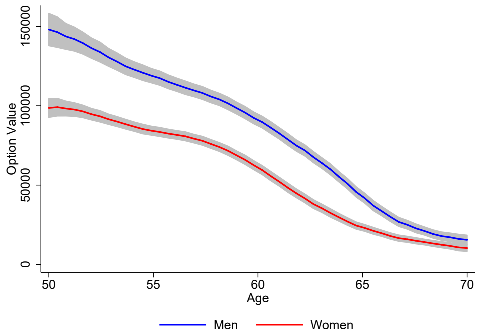

(c) Option value (DB)

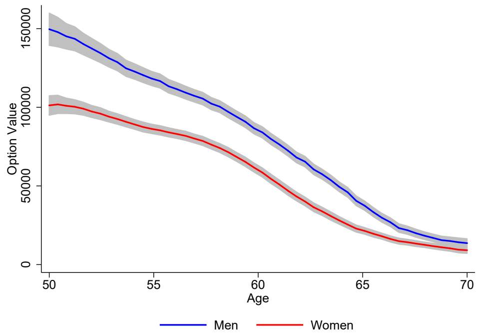

(d) Option value (DC)

Figure 3.3: Retirement Incentives by Age

ELSA. In my analysis, I control for a cubic polynomial in age, sex, the individual's total non-pension wealth (specifically their position in the non-pension wealth distribution), and total pension wealth held in state and private pensions.

5 Results

In this section I document three main sets of results. First, in section 5.1, I show that the incentives to continue working provided by private pensions are indeed associated with increased labour supply. These results update existing evidence on the labour supply effects of pensions and, to the best of my knowledge, provide the first estimates of the effects of peak value on labour supply in England.

Then, in section 5.2, I provide evidence that the labour supply response to a similarly-sized incentive to continue working are markedly higher for DB pensions than for DC. This is consistent with the labour supply incentives being more salient for DB pensions—which typically determine pension income as a relatively simple function of scheme tenure and end-of-career salary—than for DC pensions.

Finally, in section 5.3, I estimate the size of the labour supply response to pension incentives as a function of the number of distinct pensions an individual holds—a measure of how complex their pension arrangements are. I show that, for both DB and for DC pensions, labour supply responses are smaller for individuals who hold more distinct pensions.

5.1 Labour supply responses to retirement incentives

I begin by estimating the effect of financial incentives provided by private pensions on the decision to work. I estimate regression equations of the form

$$y_{it} = \alpha + \beta R_{it} + \gamma X_{it} + \epsilon_{it}, \tag{3.8}$$

where y_{it} is an indicator for whether individual i is in work in year t; $R_{it} \in \{PV, OV\}$ is a measure of the incentive to continue working provided by their private pension (either the peak value or option value); and X_{it} is a set of controls including sex, time fixed effects, a cubic in age, total wealth held in private and state pensions, the individual's non-pension wealth quintile. I estimate equation (3.8) with and without allowing for a person-specific fixed effect u_i. Including a fixed effect restricts the variation used to identify labour supply responses to changes in the marginal incentive to work within-person over time.

I present the results in Table 3.1. Both PV and OV have a positive and statistically significant effect on individuals' decision to continue working across specifications. Column (1) shows that increasing the option value of continuing to work by 1000 utils increases the probability that an individual remains in work by a statistically significant 0.14 percentage points (s.e. 0.0067 percentage points). And column (3) shows that this estimate is robust to including an individual-specific fixed effect—and so using only within-person variation in the pension incentives over time to identify labour supply responses.

Table 3.1: Labour Supply Responses to Private Pension Incentives

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	OLS	FE	FE	OLS	OLS	FE	FE
All private pension OV	0.00141***	0.00141***	0.00143***	0.00142***				
	(0.0000671)	(0.0000664)	(0.000101)	(0.0000985)				
All private pension PV					0.000513***	0.000525***	0.000209**	0.000220**
					(0.000109)	(0.000109)	(0.0000821)	(0.0000853)
Above state pension age		-0.173***		-0.155***		-0.171***		-0.163***
		(0.0111)		(0.0110)		(0.0119)		(0.0116)
Age cubic	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pension wealth	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wealth quintiles	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	26316	26316	26316	26316	26316	26316	26316	26316

Notes: Standard errors, clustered at level of the individual, in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Columns (5) and (7) show the estimated effects of the peak value measure. The estimated effects are smaller, but remain positive and statistically significant.[17] Column (5) shows that a £1000 increase in the peak value measure is associated with a 0.051 (s.e. 0.010) percentage point increase in the probability an individual works. This estimate falls to 0.021 (s.e. 0.0082) percentage points when I include an individual fixed effect in column (7), but remains statistically significant at the 5 percent level.

In even-numbered columns, I also include an indicator variable for whether the individual is above state pension age. I note two features of the results. First, the point estimates on the effects of pension incentives are little affected by the inclusion of the state pension age indicator. This reflects that, by design, the measures of pension incentive I study are unrelated to eligibility for state pensions. Second, the effect of reaching state pension age on an individual's decision to continue working is substantial. Reaching state pension age is associated with a reduction in the probability of working of between 15.5 and 17.3 percentage points across specifications. This is over and above the effect age directly, and also in addition to the effects of financial incentives provided by private pensions.

Finding a direct effect of the state pension age on retirement timing in England, despite there being very little financial incentive to actually stop work at the state pension age, is consistent with Cribb et al. (2016) and MacCuish (2019). The results in Table 3.1 demonstrate that the effect is also present when also controlling for forward-looking measures of the financial incentive to retire—and is large compared to the responses to financial incentives. It is consistent with individuals focusing attention on highly salient features of the pension system (such as the state pension age), even if they provide only limited incentive for labour supply, to the exclusion of less salient (but more substantial) financial incentives from private pensions.

Overall, this section establishes that both PV and OV measures of retirement incentive have positive and statistically significant impacts on whether an individual continues to work. It therefore updates estimates on the effect of option value on retirement timing for England, as studied previously by Banks et al. (2016). And, to the best of my knowledge, it provides the first estimates of the effect of peak value on retirement timing.

[17] I note that these estimates being smaller reflects, in part, that the two measures have different units: option value is in utils and peak value in pounds. The peak value measure also excludes labour earnings from the definition—if part of the response in option value models is driven by endogeneity arising from unoberved differences in tastes for work, the option value model would over-state the response to pension incentives.

5.2 Differences in labour supply response between pension type

I now assess whether the labour supply responses differ by the type of pension the individual holds. If labour supply responses to similarly-sized retirement incentives differ between types of pension, this would be consistent with differential awareness of the incentives. I augment regression equation (3.8) to allow for different responses to pension incentives across pension types $k \in \{DB, DC\}$,

$$y_{it} = \alpha + \sum_{k} \beta_k R_{it} + \gamma X_{it} + \epsilon_{it}. \tag{3.9}$$

In this analysis, I focus on the peak value measures of retirement incentive. There are two main reasons for this. First, the peak value measure is linear in the present value of pension wealth. This makes it possible to straightforwardly decompose the total labour supply effects documented in section 5.1 into the contributions of DB and DC pensions. Second, by excluding labour earnings it is less affected by the potential endogeneity discussed in Section 3.

I present the results in Table 3.2, estimated without and with an individual fixed effect in columns (1) and (2). In both specifications, the labour supply response to a £1000 increase in the peak value measure is stronger if that increase comes from a DB pension than a similarly-sized incentive from a DC pension. In the specification without the individual fixed effect in column (1), a £1000 increase in the peak value of a DB pension increases the probability the individual is in work by 0.090 percentage points (s.e. 0.033), compared with 0.044 percentage points (s.e. 0.010) for a DC pension. However, it is not possible to reject the hypothesis that these responses are equal at conventional significance levels: the F-statistic from a test that the coefficients are equal is 1.799, with associated p-value of 0.180.

The differences in the estimated responses between types of pension is higher for the model including an individual specific fixed effect. A £1000 increase in the peak value of a DB pension increases the probability the individual is in work by 0.077 percentage points (s.e. 0.035), compared with 0.011 percentage points (s.e. 0.0063) for a DC pension. And, importantly, in this model the hypothesis that the coefficients are equal is rejected at the 10 percent level.

The results in this section therefore provide evidence of different responses to similarly-sized financial incentives provided by private pensions, depending on whether those incentives come from a DB or a DC pension. In particular, the labour supply response to incentives provided by a DB pension is larger than the response to the same-sized incentive provided by a DC pension. This is consistent with the labour supply incentives being more salient for DB pensions—which

Table 3.2: Peak Value Models by Pension Type

	(1) OLS	(2) FE
DB peak value	0.000902***	0.000765**
	(0.000330)	(0.000352)
DC peak value	0.000444***	0.000108*
	(0.000105)	(0.0000637)
Above state pension age	-0.172***	-0.165***
	(0.0118)	(0.0115)
Age cubic	Yes	Yes
Pension wealth	Yes	Yes
Wealth quintiles	Yes	Yes
Wave dummies	Yes	Yes
F-stat (DB=DC)	1.799	3.394
p-value (DB=DC)	0.180	0.0655
Observations	26316	26316

Notes: Standard errors, clustered at level of the individual, in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

typically determine pension income as a relatively simple function of scheme tenure and end-of-career salary—than for DC pensions.

5.3 Attenuated responses if hold more pensions

The previous subsection provided evidence of labour supply responses to financial incentives differing between DB and DC pensions. This is consistent with different salience of the work incentives across the types of pension—individuals may respond more strongly to incentives provided by a DB pension because they are more salient.

In this section, I relate the differences in labour supply response to a measure of complexity of an individual's pension arrangements: the number of separate pensions they hold (either as active pensions currently being contributed to or as pensions with retained rights). If individuals are fully aware of their pension arrangements, the number of pensions they hold should have no bearing on their labour supply response to a given financial incentive. But if individuals with more complex pension arrangements are less well informed about the financial incentives provided by them to continue work, the labour supply response will fall in the number of pensions an individual holds.

In Figure 3.4 I show the distribution of the number of private pensions held across individuals in my sample. The median sample member holds two private pensions, and 75 percent of the sample hold three or fewer. However, there is a long right tail of the distribution, with over 10 percent of the sample holding five or more distinct pensions.

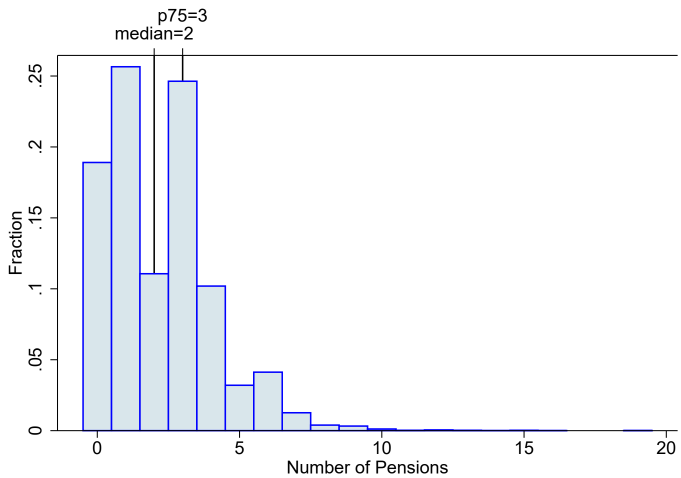

Figure 3.4: Distribution of Number of Private Pensions Held

I investigate the effect of having more complex pension arrangements on labour supply responses by further augmenting the regression equation with interaction terms between the peak value provided by pension k and the number of pensions individual i holds at time t, NP_{it}:

$$y_{it} = \alpha + \sum_k \beta_k R_{it} + \sum_k \delta_k (R_{it} \times NP_{it}) + \gamma X_{it} + \epsilon_{it}. \tag{3.10}$$

I present estimates from this specification in Table 3.3. In columns (1) and (3) I combine all private pensions together (as in section 5.1), and in columns (2) and (4) I estimate separate effects for DB and DC pensions (as in section 5.2). I include an individual-specific fixed effect in both specifications.

Across all specifications, the estimated labour supply response to private pension retirement incentives is lower for individuals with more distinct pensions. This is true for all private pensions grouped together, as well as for each of DB and DC pensions separately. This finding is consistent with individuals with more complex pension arrangements, as measured by holding more distinct pensions, being less aware of the specific financial incentives provided by their pension portfolio.

Table 3.3: Attenuated Responses if Hold More Pensions

	(1) OLS	(2) OLS	(3) FE	(4) FE
All private pension PV	0.000505***		0.000275*	
	(0.000134)		(0.000144)	
All private pens PV × no. pensions	-0.0000742***		-0.0000769**	
	(0.0000223)		(0.0000331)	
DB peak value		0.00149***		0.00196***
		(0.000409)		(0.000660)
DB PV × no. pensions		-0.000337***		-0.000429***
		(0.0000743)		(0.000124)
DC peak value		0.000507***		0.000207**
		(0.000172)		(0.000105)
DC PV × no. pensions		-0.0000792***		-0.0000797***
		(0.0000265)		(0.0000308)
Above state pension age	-0.154***	-0.155***	-0.138***	-0.140***
	(0.0116)	(0.0116)	(0.0114)	(0.0113)
Age cubic	Yes	Yes	Yes	Yes
Pension wealth	Yes	Yes	Yes	Yes
Wealth quintiles	Yes	Yes	Yes	Yes
Wave dummies	Yes	Yes	Yes	Yes
Observations	26316	26316	26316	26316

Notes: Standard errors, clustered at level of the individual, in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

6 Conclusion

This chapter has presented new estimates of the labour supply responses of older individuals in England to financial incentives provided by private pensions, and documented heterogeneity in the response to similarly sized financial incentives between DB and DC pensions. I find that the labour supply response to a £1000 accrual in pension wealth is stronger if delivered by a DB pension than a DC pension. This is consistent with the simpler rules governing the value of a DB pension being more salient to workers than the factors determining the value of a DC pension. The finding has particular policy relevance given the trend away from DB and into DC pensions in the UK. I also provide suggestive evidence that individuals' understanding of the incentives provided by their pension is influenced by the complexity of their pension portfolio—labour supply responses to a given incentive are smaller for those holding a higher number of pensions.

There are a number of potentially fruitful directions for future research. Due to data availability, I have focused on differential responses between DB and DC pensions in this chapter, but have been unable to estimate differences in response between different pensions of the same type. Testing whether individuals focus attention on a specific pension (such as the one with highest value), or certain features of each pension they hold (such as the employee rather than employer contribution) would be useful contributions to our understanding of the sources of differential salience.

Second, chapter has focused solely on the incentives provided by private pensions on the extensive margin around retirement age. However, pensions provide labour supply incentives to individuals throughout their working lives. For individuals contributing to a DB pension, an additional year in the scheme may increase the fraction of late-career earnings paid during retirement. Similarly, employers typically contribute some fraction of employees' earnings to their DC pension. Both of these features influence the effective wage that an individual earns in each year of their working life—although the component of pay operating through the pension is delayed. It would be informative to test whether variation in the effective wage operating through pension incentives affects labour supply similarly to variation induced by, say, the tax and transfer system. Conducting tests of this type would likely benefit from a structural model of life-cycle labour supply.

Finally, the evidence in this chapter only documents differential responses to similarly-sized incentives. While it is likely that these differences arise from under-response to non-salient pension features, it is not possible to rule out that the differences may arise from over-responses to highly salient features. More generally, it would be useful to study the individual effects of different pension features on retirement, including financial features (such as accrual rates) and non-financial features (such as retirement ages), to assess the extent to which each is under- or over-weighted. Again, developing a structural model of labour supply in old age would provide a benchmark against which to compare the empirical responses, and so to assess whether they are higher or lower than would be implied by a model with perfect salience.

Bibliography

Abeler, J. and Jäger, S. (2015). Complex Tax Incentives. *American Economic Journal: Economic Policy*, 7(3):1–28.

Agarwal, S., Driscoll, J. C., Gabaix, X., and Laibson, D. (2008). Learning in the credit card market. Technical report, National Bureau of Economic Research.

Attanasio, O. and Weber, G. (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy*, 103(6):1121–57.

Banks, J., Emmerson, C., and Tetlow, G. (2016). Effect of Pensions and Disability Benefits on Retirement in the United Kingdom. In Wise, D. A., editor, *Social Security Programs and Retirement Around the World: Disability Insurance Programs and Retirement*, pages 81–136. University of Chicago Press.

Blanden, J. and Gregg, P. (2004). Family Income and Educational Attainment: A Review of Approaches and Evidence for Britain. *Oxford Review of Economic Policy*, 20(2):245–263.

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

Blundell, R., Costa Dias, M., Meghir, C., and Shaw, J. (2016a). Female labor supply, human capital, and welfare reform. *Econometrica*, 84(5):1705–1753.

Blundell, R., Duncan, A., and Meghir, C. (1998). Estimating labor supply responses using tax reforms. *Econometrica*, pages 827–861.

Blundell, R., French, E., and Tetlow, G. (2016b). Retirement incentives and labor supply. *Handbook of the Economics of Population Aging*, 1:457–566.

Blundell, R. and MaCurdy, T. (1999). Labor supply: A review of alternative approaches. In *Handbook of labor economics*, volume 3, pages 1559–1695. Elsevier.

Blundell, R., Meghir, C., and Smith, S. (2002). Pension incentives and the pattern of early retirement. *The Economic Journal*, 112(478):C153–C170.

Brewer, M. (2003). The new tax credits. Technical Report 35.

Carneiro, P. M. and Heckman, J. J. (2003). Human Capital Policy. SSRN Scholarly Paper ID 434544, Social Science Research Network, Rochester, NY.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and Taxation: Theory and Evidence. *American Economic Review*, 99(4):1145–77.

Chevalier, A. and Lanot, G. (2002). The Relative Effect of Family Characteristics and Financial Situation on Educational Achievement. *Education Economics*, 10(2):165–181.

Coile, C. and Gruber, J. (2007). Future Social Security Entitlements and the Retirement Decision. *The Review of Economics and Statistics*, 89(2):234–246.

Crawford, R. (2012). ELSA Pension Wealth Derived Variables (Waves 2 to 5): Methodology. Technical Report 1.

Crawford, R. and Tetlow, G. (2010). Employment, retirement and pensions. In Banks, J., Lessof, C., Nazroo, J., Rogers, N., Stafford, M., and Steptoe, A., editors, *Financial circumstances, health and well-being of the older population in England: The 2008 English Longitudinal Study of Ageing (Wave 4)*, page 11.

Cribb, J., Emmerson, C., and Tetlow, G. (2016). Signals matter? Large retirement responses to limited financial incentives. *Labour Economics*, 42:203–212.

Dalton, C. M., Gowrisankaran, G., and Town, R. J. (2019). Salience, Myopia, and Complex Dynamic Incentives: Evidence from Medicare Part D. *The Review of Economic Studies*.

Department for Work and Pensions (2015). 2010 to 2015 government policy: welfare reform.

Dolls, M., Doerrenberg, P., Peichl, A., and Stichnoth, H. (2018). Do retirement savings increase in response to information about retirement and expected pensions? *Journal of Public Economics*, 158:168–179.

Feldman, N., Goldin, J., and Homonoff, T. (2018). Raising the Stakes: Experimental Evidence on the Endogeneity of Taxpayer Mistakes. *National Tax Journal*, 71(2):201–230.

Feldman, N. E., Katuščák, P., and Kawano, L. (2016). Taxpayer Confusion: Evidence from the Child Tax Credit. *American Economic Review*, 106(3):807–835.

French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *The Review of Economic Studies*, 72(2):395–427.

Grubb, M. D. and Osborne, M. (2015). Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock. *American Economic Review*, 105(1):234–271.

Gustman, A. L. and Steinmeier, T. L. (2004). Social security, pensions and retirement behaviour within the family. *Journal of Applied Econometrics*, 19(6):723–737. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/jae.753.

HMRC (2017). Child and Working Tax Credits Statistics: Final annual awards in 2012-13.

Ito, K. (2012). Do consumers respond to marginal or average price? Evidence from nonlinear electricity pricing. Technical report, National Bureau of Economic Research.

Joint Committee on Taxation (2015). Complexity in the Federal Tax System (JCX-49-15).

Ketcham, J. D., Lucarelli, C., Miravete, E. J., and Roebuck, M. C. (2012). Sinking, Swimming, or Learning to Swim in Medicare Part D. *The American Economic Review*, 102(6):2639–2673.

Kim, D. (2020). Worker retirement responses to pension incentives: Do they respond to pension wealth? *Journal of Economic Behavior & Organization*, 173:365–385.

Laibson, D. (1997). Golden eggs and hyperbolic discounting. *The Quarterly Journal of Economics*, 112(2):443–478.

Liebman, J. B. and Luttmer, E. F. (2015). Would people behave differently if they better understood social security? Evidence from a field experiment. *American Economic Journal: Economic Policy*, 7(1):275–99.

Liebman, J. B. and Zeckhauser, J., R. (2004). Schmeduling. *NBER Working Paper*.

Lusardi, A. and Mitchell, O. S. (2007). Baby Boomer retirement security: The roles of planning, financial literacy, and housing wealth. *Journal of Monetary Economics*, 54(1):205–224.

MacCuish, J. H. (2019). Rational Inattention and Retirement Puzzles. *Available at SSRN 3363188*.

MaCurdy, T. E. (1981). An empirical model of labor supply in a life-cycle setting. *Journal of political Economy*, 89(6):1059–1085. Publisher: The University of Chicago Press.

Mastrobuoni, G. (2011). The role of information for retirement behavior: Evidence based on the stepwise introduction of the Social Security Statement. *Journal of Public Economics*, 95(7):913–925.

McAlpine, C. and Thomas, A. (2008). The triggers and barriers to the take-up of Working Tax Credit among those without dependent children. *HM Revenue and Customs Research Report*, 86:72.

McClelland, R. and Mok, S. (2012). A review of recent research on labor supply elasticities.

Meghir, C. and Phillips, D. (2010). Labour supply and taxes. *Dimensions of tax design: The Mirrlees review*, pages 202–74.

Mirrlees, J., Adam, S., Besley, T., Blundell, R., Bond, S., Chote, R., Gammie, M., Johnson, P., Myles, G., and Poterba, J. M. (2011). Integrating Personal Taxes and Benefits. In *Tax by Design*. Oxford University Press.

Rohwedder, S. and Kleinjans, K. J. (2006). Dynamics of Individual Information about Social Security. *Unpublished manuscript, RAND*. Publisher: Citeseer.

Saez, E. (2010). Do taxpayers bunch at kink points? *American economic Journal: economic policy*, 2(3):180–212.

Seibold, A. (2019). Reference Points for Retirement Behavior: Evidence from German Pension Discontinuities. Publisher: CESifo Working Paper.

Smith, J. C. (2011). The Ins and Outs of UK Unemployment*. *The Economic Journal*, 121(552):402–444.

Stango, V. and Zinman, J. (2014). Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Bank Overdraft Fees. *The Review of Financial Studies*, 27(4):990–1030.

Stock, J. H. and Wise, D. A. (1990). Pensions, the Option Value of Work, and Retirement. *Econometrica*, 58(5):1151–1180. Publisher: [Wiley, Econometric Society].

Taubinsky, D. and Rees-Jones, A. (2017). Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment. *The Review of Economic Studies*.

Wingender, P. and LaLumia, S. (2017). Income Effects on Maternal Labor Supply: Evidence from Child-Related Tax Benefits. *National Tax Journal*, 70(1):11–52.