Benefit Salience and Labour Supply

Peter Spittal

Discussion Paper 22/764

5 April 2022

School of Economics

University of Bristol
Priory Road Complex
Bristol
BS8 1TU
United Kingdom

University of BRISTOL
School of Economics

Benefit Salience and Labour Supply

Peter Spittal*

April 2022

Abstract

I study the salience of dynamic incentives provided by the welfare system, as revealed by labour supply responses to foreseeable reductions in benefit income. I show that claimants fail to anticipate a large lump-sum reduction in benefit entitlement, arising predictably from children ageing out of eligibility for the UK's Child Tax Credit. I show also that the salience of the rules increases with experience. I then develop a structural life-cycle labour supply model incorporating potential non-salience of eligibility rules. The model estimates suggest that 82 percent of claimants initially fail to anticipate the benefit reduction. The resulting optimisation errors have substantial welfare costs—equivalent to a 14 percent reduction in income from the programme, with no offsetting benefits to the government. The findings reveal a previously undocumented source of inefficiency in the welfare system, arising from non-salient policy features with significant financial consequences.

JEL classification: D15, D83, H31, J22.

*School of Economics, University of Bristol. Email: peter.spittal@bristol.ac.uk. I am grateful to Eric French for valuable advice and guidance, as well as to Richard Blundell, Monica Costa Dias, Jan Eeckhout, Amir Habibi, Jamie Hentall-MacCuish, Paul Hufe, Gregory Jolivet, Attila Lindner, Rory McGee, Matthais Parey, Suphanit Piyapromdee, Fabien Postel-Vinay, Uta Schönberg, Hans Sievertsen and seminar participants at Bath, Bristol, Exeter, the Institute for Fiscal Studies, Nottingham, Royal Holloway, Reading, Surrey, University College London and Université catholique de Louvain for helpful comments and discussions. I gratefully acknowledge financial support during this project from an ESRC doctoral studentship.

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 not anticipate the full financial implications of their decisions or changes in their circumstances. In particular, people may underweight or ignore policy features which are not *salient*. Such effects have been documented in the context of sales taxes (e.g. Chetty et al., 2009; Feldman et al., 2018; Taubinsky and Rees-Jones, 2017), income taxes (e.g. Saez, 2010) and income-related benefits (e.g. Chetty and Saez, 2013). This has important implications for both designing and assessing policy.

In addition to the static policy features typically studied in existing literature, many public policies have important *dynamic* elements. For example, it is common for benefit entitlement to vary deterministically over time or with claimant characteristics such as age or household composition. The financial implications of dynamic policy features can be substantial, and so failing to account for them fully could lead to sizeable optimisation errors and, consequently, welfare costs.

In this paper, I study the salience of a major dynamic feature of the UK welfare system. Like many countries, the UK provides additional benefit payments to households with dependent children; children ageing out of eligibility for such payments therefore induces a large, exogenous, lump-sum and foreseeable reduction in their parents' benefit entitlement. The central question is whether these child-related eligibility rules are salient for claimant families and, if not, what are the welfare costs? I show that the rules are non-salient, despite the high financial stakes, but that the salience increases through experience: families who have already experienced a benefit reduction better anticipate future reductions. The welfare costs are substantial—equivalent to a 14 percent reduction in the generosity of the benefit payment, with no offsetting benefits to the government. This points to a previously undocumented source of inefficiency in the welfare system.

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, a standard life-cycle labour supply model would predict no contemporaneous labour supply response: instead, claimants would work more in advance of the reduction, accumulating savings to better smooth marginal utility at the time the benefits are withdrawn. The existence of a causal contemporaneous labour supply response sheds light on the extent to which claimants anticipated the benefit reduction, and hence whether they were aware of the eligibility rules. Particular features of the policy allow me also to investigate directly whether the salience of the rules increases with 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 the UK's Child Tax Credit (CTC), which was the main source of child-related support for low-income families for the 10 years from 2003 until the start of its phased replacement in 2013. Over this period, it reached an average of 5.5 million

[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.

families each year at an annual cost of £17.5 billion (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 benefit income.

I use a difference-in-differences empirical strategy to identify the causal effect of reductions in CTC on labour supply. Using monthly panel data on the labour supply and family characteristics of cohabiting mothers, drawn from the UK Household Longitudinal Study (UKHLS) linked to its predecessor the British Household Panel Survey (BHPS), I compare the labour supply of 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] This allows me to identify any labour supply response to reduced benefit entitlement separately from any direct effect the child's age could have on their mother's labour supply.

I begin by estimating the average labour income 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 contemporaneous labour income response to a lump-sum and foreseeable benefit reduction is inconsistent with the predictions of a standard life-cycle labour supply model. But it is consistent with the age-related eligibility rules having been non-salient.

I then document two further results which are consistent with the salience of the eligibility rules increasing through experience, and also help to rule out a number of alternative mechanisms. 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.

First, I show that there is only a labour income response among mothers who experience a CTC reduction for the first time. This is despite that, due to the way CTC is structured, the reduction in benefits is at least as large as for the first child when subsequent children become ineligible. 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 the salience of the rules increasing with experience.

Second, I estimate separate labour supply effects for families with different numbers of still-eligible children. All families experience the same *contemporaneous* reduction in CTC when one of their children becomes ineligible, but *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 for otherwise forward-looking households. If it does, the labour supply response will be larger for families with

[2]Despite being called a "tax credit", CTC was a transfer paid into claimants' bank accounts each month and was 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 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: in particular (1) the labour supply of mothers in the months leading up to the child becoming ineligible follow a common trend and (2) there are no statistically significant differences in household structure at the time the household becomes ineligible which could explain differential labour supply behaviour.

more still-eligible children, as increased salience of the eligibility rules generates a larger shock to the present value of benefit income. I find that this is indeed the case.[4]

Taken together, these results provide new empirical evidence that a *dynamic* feature of a major welfare policy with high financial stakes is imperfectly salient, and that the salience of the policy rules increases with experience. The findings are also inconsistent with a number of other explanations, including: liquidity constraints, time inconsistency, purely myopic households, discreteness in the labour supply decision, or mistaken understanding of the reason for the benefit reduction. Identifying the effects of an initially non-salient policy feature separately from these alternative explanations, which are indistinguishable in other settings but have different implications for policy and welfare, is a benefit of the dynamic environment I study and a key contribution of this paper.

However, a number of important questions cannot be answered directly from the reduced form results. Specifically, how widespread is the non-salience of the eligibility rules? 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 the age-related eligibility rules may not be salient. I estimate key parameters of the model by indirect inference, using the empirical results as targets. In particular, in conjunction with the structure of the model, the empirical results provide sufficient variation to identify separately (i) preference parameters which govern labour supply responsiveness and (ii) the proportion of households for which the eligibility rules are non-salient.

The estimated preference parameters are closely in line with existing literature, underlining that the reduced form results are also consistent quantitatively (as well as simply qualitatively) with introducing non-salient eligibility rules to an otherwise-standard lifecycle labour supply model. I also estimate that 82 percent of claimants are unaware of the benefit rules at the time their first child becomes ineligible, suggesting widespread failure of households to fully anticipate the impact of children ageing out of eligibility on their household budget.

Finally, I assess the consequences for welfare. I calculate the reduction in initial assets households would be willing to accept to be aware of the eligibility rules over their entire lifecycle. This compensating variation is substantial—equal to 14 percent of the lifetime present value of CTC. The large welfare cost arises for two reasons. First, the financial impact of losing eligibility for CTC is itself large. Second, and importantly, the existence of age-related eligibility rules provides dynamic incentives for labour supply over the entire lifecycle. Failing to anticipate the benefit reduction therefore leads to optimisation errors which compound over time, necessitating especially costly adjustment when children become ineligible. I also show that, unlike other widely-studied settings such as sales taxes (e.g. Farhi and Gabaix, 2020; Goldin, 2015), there are no offsetting benefits to the government from non-salient rules. This paper therefore identifies major source of inefficiency in the welfare system which has not previously been documented.[5]

[4] In principle, the asymmetric response between first and subsequent children could be driven by discreteness in labour supply (meaning that households cannot adjust labour supply repeatedly) rather than an increase in the salience of the eligibility rules. However, the asymmetry in responses between families with different numbers of still-eligible children could not be explained by discreteness in labour supply alone, without there also being an increase in the salience of the eligibility rules in forward-looking households. I discuss this point further in Section 4.4.

[5] In fact, because I abstract from labour supply frictions in the structural model, which would introduce costs

Contribution. This paper makes three 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 non-salient and leads claimants to make optimisation errors. Existing empirical literature on salience in tax systems has typically focused on static environments which are arguably simpler and often have substantially lower financial stakes, such as sales taxes (e.g. Chetty et al., 2009; Feldman et al., 2018; Taubinsky and Rees-Jones, 2017).[6] It is notable to find evidence that features of a welfare system with much higher financial stakes, and which provide dynamic incentives for labour supply and consumption, may also not be salient. Documenting the consequences for claimants' labour supply decisions over their life-cycle highlights a new source of inefficiency in the welfare system.

Second, I provide evidence that the salience of policy features may increase directly with 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 paper to provide direct evidence that the salience of features of the welfare system increases following unexpected financial consequences.

This has a number of important implications. It suggests that the frequency at which people interact with a system may determine how salient its features are: even sizeable financial incentives can be non-salient when they are part of a system people receive feedback from infrequently. It also helps to distinguish the effects of non-salient eligibility rules from other potential frictions (such as hyperbolic discounting or myopia) which would not lessen with experience.

Third, to the best of my knowledge, this is the first paper to incorporate non-salient dynamic incentives into a structural model of life-cycle labour supply. This contributes to the literature using structural methods to investigate the consequences of behavioural frictions for public policy (e.g. Dalton et al., 2020; Fang and Wang, 2015; Paserman, 2008). It also contributes to the structural life-cycle labour supply literature (e.g. Blundell et al., 2016) by adding an empirically-relevant friction which affects the assessment of optimal policy. The structural analysis allows me both to validate the reduced form results, showing that non-salient eligibility rules are capable of generating the responses quantitatively as well as qualitatively, and to assess the welfare implications. The sizeable welfare consequences underline that considering the salience of dynamic policy features is important, as optimisation errors can accumulate over time and become especially costly.

The findings of this paper also contribute to policy debates surrounding the complexity of welfare systems (e.g. Mirrlees et al., 2011, for the UK). In the specific context of the UK, a

associated with labour supply adjustment, the welfare costs I estimate are likely to be a conservative lower bound on the true costs of the non-salient rules.

[6] An exception is Dalton et al. (2020), 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.

number of papers have documented that the 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 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).[7] This paper provides new evidence that people did systematically misunderstand important features of the welfare system, and the welfare cost was substantial.

Finally, the empirical setting of this paper is related to a number of papers using predictable changes in child-related benefits to estimate labour supply responses (e.g. Feldman et al., 2016; Wingender and LaLumia, 2017). For example, Feldman et al. (2016) study the labour supply impact of children ageing out of eligibility for the US version of the Child Tax Credit. Differently from this paper, Feldman et al. (2016) find evidence that households mistake a lump-sum reduction in benefit income with an increase in the effective *marginal* tax rate, and hence reduce labour supply. I find no evidence of households making such errors in the context of the UK CTC: indeed, the reduced form results are inconsistent with households making mistakes of this form.[8] Instead, I document that the age-related eligibility rules were not salient and, using data which allow me to identify mothers who lose eligibility repeatedly, that the optimisation errors lessen with experience.

Outline of the paper. The rest of the paper is structured as follows. I provide further details of CTC in Section 2. I then discuss my empirical strategy for identifying the pattern of labour supply responses, and the data I use, in Section 3. Section 4 sets out my main results and Section 5 contains a number of robustness tests. I then set out a life-cycle labour supply model in which the rules governing eligibility for certain benefits are potentially not salient in Section 6. I use the pattern of labour supply responses to estimate key parameters of the model, and calculate the welfare implications of the non-salient eligibility rules. Section 7 concludes.

2 Details of Child Tax Credit

Child Tax Credit was 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 to low-income parents known as Working Families' Tax Credit (Brewer, 2003). Since April 2013, tax credits have been replaced for some claimants by the phased roll-out of a new benefit known as Universal Credit.[9] In this paper, I focus on the 10 year period from April 2003 to April 2013 in which CTC was available to all claimants.

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

[8] Despite having the same name, the UK and US CTCs are very different schemes. This likely accounts for differences in the findings. For example, the US CTC is usually administered as a tax refund, reducing a claimant household's tax liability for the year. By contrast, CTC in the UK is paid as a monthly transfer, unrelated to the household's tax liability.

[9] 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.

Table 1: Child Tax Credit Parameters

	Maximum entitlement		Taper thresholds and rates			
	Family element ϕ_f	Child element ϕ_c	First income threshold $\bar{Y}_1$	First taper rate $\bar{\tau}_1$	Second income threshold $\bar{Y}_2$	Second taper rate $\bar{\tau}_2$
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.

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 age 20 if they remain in full-time non-advanced education.[10] 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.[11] 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 dependent children.[12] A family's maximum entitlement is therefore given by

$$\bar{b}(k) = \mathbb{1}(k > 0)(\phi_f + \phi_c k).$$

For example, in the 2012/13 tax year, the family element ϕ_f was £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.

A household's actual CTC entitlement is also determined by its 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,

[10] 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.

[11] In Online Appendix A.2, I provide direct 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, using an instrumental variables strategy to isolate exogenous variation in eligibility arising only from child ageing.

[12] 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.

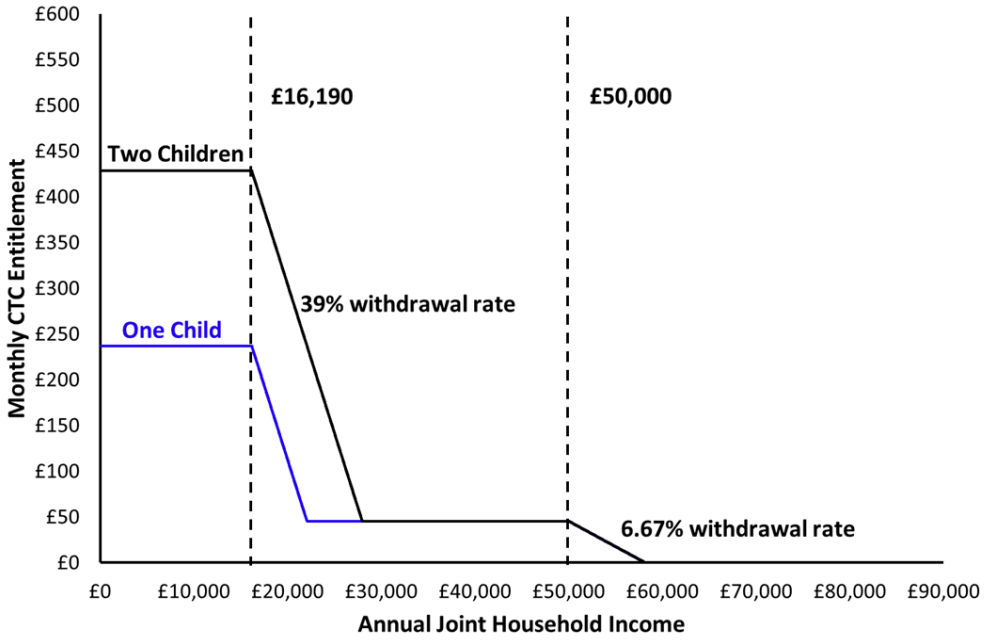

Figure 1: CTC schedule in 2010/11 tax year

the family element was withdrawn at rate 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.[13] 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 shows the parameters determining CTC entitlement in each tax year from 2003/04 until 2012/13, and Figure 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 fraction of a family's claim which has been withdrawn as $\tau(k, Y)$. A family's actual entitlement to CTC is therefore given by

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

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 households 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 (and so experience a decrease in their effective marginal tax rate); while others still will experience no change at all because they are on a part of the benefit schedule which

[13] 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.

is invariant to k. This variation across households 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 (i) any substitution effect arising from a change in the withdrawal rate or (ii) 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 under the UK CTC 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 families 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 my empirical approach is to estimate labour supply responses only for those households for which losing benefit eligibility has no effect on the effective marginal tax rate: households for which the effective marginal tax rate also falls when benefits are removed may increase labour supply due to a substitution effect, even if the rules are fully salient.

3 Empirical Strategy and Data

3.1 Empirical Strategy

The focus of this paper is whether labour supply responds contemporaneously after a foreseeable lump-sum reduction in benefits, arising from children ageing out of eligibility for the UK's CTC. A standard life-cycle labour supply model would predict no contemporaneous labour supply response in this context: instead, claimants would work more in advance of the reduction, accumulating savings to better smooth marginal utility at the time the benefits are withdrawn. The existence of a causal contemporaneous labour supply response sheds light on the extent to

which claimants anticipated the benefit reduction, and hence whether they were aware of the eligibility rules.

The main empirical challenge is to identify the response to reduced benefits separately from any *direct* effects child ageing may have on maternal labour supply. I use a difference in differences empirical strategy which compares the labour supply of mothers who experience a lump-sum reduction in benefit income when their child becomes ineligible for CTC against a control group with children of the same age who experience no change in CTC entitlement. As all mothers in the analysis have children of the same age, this holds constant the direct effect of child age on labour supply.

Treatment groups. 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 (rather than any changes in eligibility which arise due to subsequent labour supply adjustments).

There are three groups. Group C is the control group of households which experience no change in CTC when their child becomes ineligible. These households' combined household income puts them on a section of the CTC schedule where the actual and counterfactual entitlements are identical in every month before the child becomes ineligible (i.e. they are entitled only to the family element or are entirely ineligible to claim). Requiring that invariance of the benefit entitlement in *each* of the 24 months before the child becomes ineligible is conservative: if a household's benefit entitlement would have been different in even one of the 24 months before the child becomes ineligible, I consider it to be treated.

There are two treatment groups. Group L contains households which experience only a lump sum reduction in CTC when the child becomes ineligible. These households have a counterfactual benefit entitlement below their actual entitlement for at least one of the 24 months before the child becomes ineligible, but identical actual and counterfactual withdrawal rates in each month. The behaviour of this group is the main focus of the paper as it sheds light on whether they anticipated the benefit reduction: there is no substitution effect, and so a labour supply response is consistent only with having failed to have anticipate the loss of benefits. Again, the criteria for assigning to group L are conservative: excluding households for which the actual and counterfactual withdrawal rates would have differed in any one of the 24 months before losing eligibility ensures, as far as possible, that group L members do experience only a lump-sum benefit reduction.

Finally, group M contains households for which a child ageing out of eligibility leads to a reduction in both the level of benefits and the withdrawal rate. This group contains households 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 paper is to test for the existence of an income effect, for which group L is sufficient; I estimate responses for group M only to separate any response of this group from L.

Difference-in-differences specification. I estimate labour supply responses for mothers in each treatment group by estimating equations of the form

$$y_{it} = \alpha + \sum_{g \in \{L,M\}} \beta_g (T_i^g \times A_{it}) + \sum_{g \in \{L,M\}} \gamma_g T_i^g + \delta A_{it} + \zeta X_{it} + \theta_t + \epsilon_{it}, \tag{2}$$

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 $\{\beta_g\}_{g \in \{L,M\}}$ are the labour supply responses in each treatment group. 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 $\beta_L > 0$.

Interpreting estimates of β_g as the causal effects of CTC changes for treatment group g requires that (i) 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 (ii) 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 5.[14]

Additional tests. There are two further testable implications of the salience of the eligibility rules increasing with experience. First, having learned that benefit receipt depends child age, there should be no income effect at the time of any subsequent children becoming ineligible. This is because the household incorporates all future reductions in CTC into its marginal utility of lifetime wealth at the time the first child becomes ineligible—and so adjusts labour supply to enable it to smooth marginal utility at the time of all future benefit reductions.

Second, the income effect on labour supply at the time of learning about the eligibility rules will be larger for mothers with more still-eligible children. These households experience a larger shock to perceived lifetime wealth when they learn of the benefit eligibility rules, leading to a larger increase in labour supply.

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 (2) which allow for separate treatment effects between these subgroups.

[14] A potential threat to condition (ii) arises because parents of children aged between 16 and 20 retain eligibility only if their child remains in full time education. In Online Appendix A.2 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.

3.2 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 each household's CTC entitlement, and the taper rate it faces, based on reported household income, hours worked and number of dependent children. I also calculate each household's counterfactual entitlement and taper rate, i.e. the amount of CTC it would have been eligible to claim, and the taper rate it would have faced, if it had one fewer dependent child. I provide further details on how I constructed the data in Online Appendix B.

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). My strategy is to study labour supply responses as a means of inferring the extent to which households anticipated the benefit reduction; it is therefore important to focus on a group whose labour supply is most likely to be responsive.[15]

I further restrict the sample to mothers who cohabit with a partner. This provides an additional source of variation in CTC, which is assessed based on the combined income of both parents—mothers who earn the same as each other, but have partners who earn different amounts, experience different reductions in CTC when their children age out of eligibility.[16] However, I show in Section 5 that the pattern and magnitude of my results are robust to also including single mothers in the sample.

Finally, I keep only households 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—and focus on the 24 months either side of losing eligibility in the main analysis.

The final sample includes 1,204 mothers with 1,870 children who lost eligibility for CTC between April 2003 and April 2013. Table 2 contains descriptive statistics for the final sample as a whole, as well as for the treatment and control groups I describe in Section 3.1.

4 Results

I establish three main results about the labour supply response of cohabiting 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, which increases 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

[15] In Online Appendix Section A.5, I show that I indeed find no effect of losing eligibility on male labour supply.

[16] 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 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	

Notes: Table shows descriptive statistics for the entire sample of cohabiting mothers, all together and separately by the three main treatment groups. See text in Section 3.1 for details on the treatment groups.

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 the salience of the eligibility rules increasing for households after experiencing the benefit reduction, and are inconsistent with the eligibility rules being fully salient. And, as I discuss in Section 4.4, the results are also inconsistent with a number of alternative models which could also generate some (but not all) features of the main results.

4.1 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 2 compares the labour supply of mothers in the control group C against those in treatment group L. Each dot is the estimated β_L coefficient from equation (2), amended to allow separate effects in each two-month period, normalised to equal zero in the two month period before losing eligibility. The grey area shows the 90 percent confidence interval.

The labour income of the groups evolves similarly in the months before losing eligibility for CTC, but begins to diverge exactly at the time treatment group members lose eligibility for CTC. The labour supply effect continues to materialise over the course of 24 months (likely reflecting frictions in labour supply adjustment), before levelling off. As labour income appears to adjust fully within two years, I restrict attention to the two years either side of losing eligibility in all subsequent analysis. Overall, Figure 2 demonstrates that there is a causal effect of the benefit reduction on the labour supply of mothers in treatment group L.

Next I estimate regression equation (2), 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 3. Panel A contains the estimated treatment effects on mothers' monthly labour income, monthly CTC receipt and the

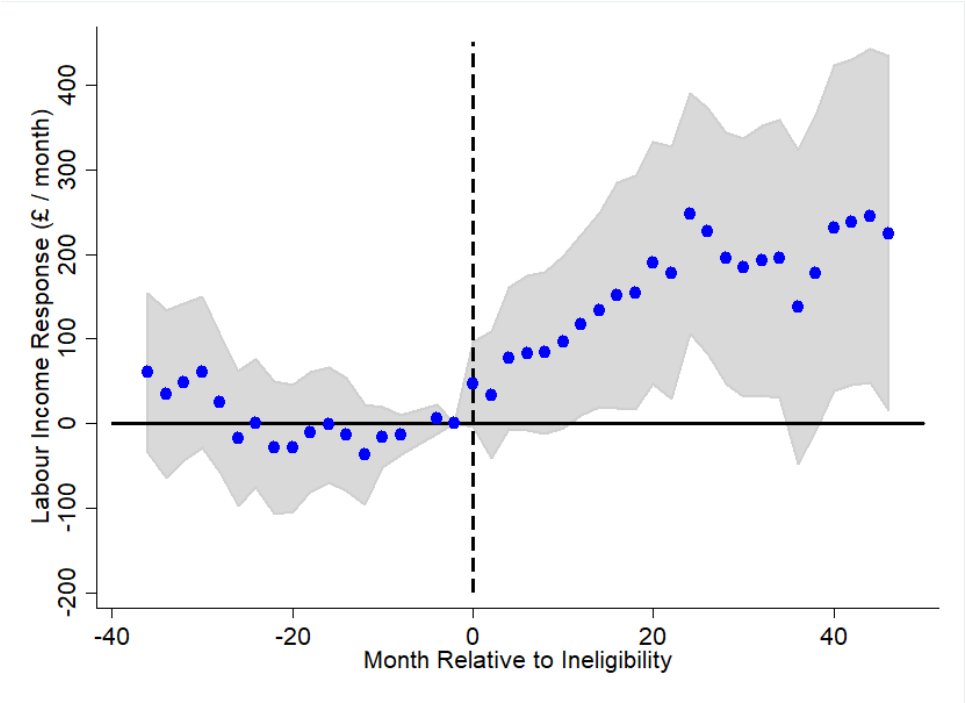

Figure 2: Labour Supply Effect of Losing CTC. Notes: Figure shows the difference in maternal labour income between treatment group L and control group C in each two-month period around their child becoming ineligible for CTC. Each dot is the estimated β_L coefficient from equation (2), amended to allow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

withdrawal rate. These are estimates of β_L from equation (2), 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.[17] 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 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 the estimated response as elasticities with respect to unearned income.[18] I distinguish two cases. The "myopic" elasticities assume that all of the estimated labour supply adjustment is a response only to the contemporaneous change in unearned income. The resulting labour income elasticity is -0.61. 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 will fall further in future.

The "forward-looking" elasticities take this into account. In these calculations, I express the estimated percentage labour supply adjustment as a percentage of 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 contemporaneous reduction.[19] This assumption reduces the labour income elasticity to -0.22.

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

[17] In Online Appendix A.1 I estimate the effect of losing eligibility on a binary measure of labour supply. The pattern of point estimates for this binary measure is similar to that discussed here, but the magnitudes are smaller: taken together with the results on labour income, this points to the existence of a large intensive margin response.

[18] I computed the elasticities by (1) calculating the percent changes in labour supply and unearned income by dividing 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 to estimate standard errors.

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

Table 3: Labour Supply Response to Lump Sum CTC Reduction

	All	First Child	Other Child	First Child, Large Family	First Child, 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)
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)
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. Column (1) shows the treatment effect estimated at the time any child ages out of eligibility discussed in Section 4.1; Columns (2) and (3) distinguish between first and subsequent children as discussed in Section 4.2; Columns (4) to (6) distinguish first children from large vs. small families as discussed in Section 4.3. For further details see the main text in Section 4. 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.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

eligibility rules through experience.

4.2 Income effect only when first child loses eligibility

A testable implication of the salience of benefit eligibility rules increasing with experience is that any income effect on labour supply should be strongest the first time mothers have a child become ineligible for CTC, as mothers would adjust labour supply to account for both the contemporaneous and any future expected changes to benefit income.

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 estimate a version of equation (2) which includes these 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.

The estimated treatment effects for group L_{first} in are in column (2) of Table 3, and the estimates for L_{other} are in column (3). The labour supply response is large and statistically significant for mothers who lose eligibility for CTC for the first time: they increase labour income by an average of £158.9 a month, (s.e. £59.0). This is larger than the total average response for 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 solely 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. The estimated effects are consistent with the age-related eligibility rules becoming fully salient after families first lose entitlement to CTC, with no further adjustments to labour income at the time of subsequent benefit reductions.[20]

Panel B expresses these estimated treatment effects as elasticities. The myopic labour income elasticity is -0.77 for the first child, while the forward-looking elasticity is mechanically smaller at -0.30. Both 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 effect 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 exceptionally high replacement rate also suggests that labour supply responds to more than just the *contemporaneous* benefit loss: instead,

[20] One possible concern is that the asymmetric labour supply response at the time of first and subsequent children ageing out of eligibility for CTC is driven by discreteness in mothers' labour supply decisions: after increasing labour supply once, there may be no scope to increase labour supply further. I discuss this point further in Section 4.4, and conclude that discreteness in labour supply alone cannot explain the full pattern of labour supply responses.

labour supply also adjusts to the increased salience of additional benefit reductions which will arise in future.

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.

4.3 Income effect larger for mothers with more dependent children

If the salience of benefit eligibility rules increases through experience, the labour supply response at the time a household's first child becomes ineligible should be larger among mothers with 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 the first child becoming ineligible is different between those with many and few remaining dependent children . 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.[21] I then divide treatment group L into three ($L_{first,large}$, $L_{first,small}$ and L_{other}) and allow the labour supply responses at the time a family's first child becomes ineligible to differ across these groups.

Columns (4) to (6) of Table 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. 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.[22] Also consistent with this, the estimated labour supply effect at the time subsequent children age out of eligibility remain numerically small and statistically insignificant.

Panel B expresses these treatment effects as elasticities. The myopic elasticity for large families are very large, with an income elasticity of -2.46; the elasticity for smaller families are numerically smaller, at -0.60 , but still large compared to estimates from existing literature (e.g. McClelland and Mok, 2012). Again, I interpret these exceptionally large contemporaneous responses as indicating that the increases in labour supply are a response also to the knowledge of future benefit reductions. The forward-looking elasticites, which assume that the labour supply responses take account of all future CTC reductions, are more in line with existing literature.

Finally, Panel C expresses the estimates as income replacement rates. The myopic replacement rates, which assume labour income responds only to the contemporaneous CTC reduction, at the time a family's first child ages out of eligibility are again very high for both large and

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

[22] In Section 6, I show that the magnitudes of these results are also consistent quantitatively with adding non-salient eligibility rules to an otherwise-standard life-cycle model of labour supply.

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. As before, the exceptionally high myopic replacement rates suggest that labour supply responds to more than just the contemporaneous loss in benefit income. This points to the increased salience of future benefit reductions also driving the labour income response.

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. 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.

4.4 Alternative mechanisms

The results set out above are consistent with households learning of the rules governing eligibility to CTC through experience. I now discuss a number of alternative mechanisms which could also be consistent with some features of the empirical results, and argue that none is consistent with *all* features. Being able to rule out these alternatives is important, as many would have different implications for policy and the welfare consequences of the findings.

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 (and work less) than a time consistent worker 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 of the benefit reduction, 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, such an explanation would not be consistent with the results presented in Sections 4.2 and 4.3. If households are time inconsistent, they would fail to accumulate sufficient savings to smooth marginal utility at the time of *all* reductions in benefits, not just the first. Finding that the labour supply response is driven solely by households losing eligibility for the first time is inconsistent with this. Equally, for a time inconsistent household, there would be no reason to increase labour supply by more, at the time of the first child becoming ineligible, if the household expected to have further benefit reductions in future: if the eligibility rules were fully salient, the household would already be aware of any future benefit reductions.

Myopia. Myopic households discount the future arbitrarily heavily, and so fail to plan for *any* future changes in their benefit entitlement. Such households are distinct from those for which the salience of the eligibility rules increases with experience because they never foresee benefit reductions associated with their children becoming ineligible for CTC, even if they have experienced them before.

If households 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: the households would not smooth marginal utility at the time of any benefit reduction. Equally, it would certainly not be the case the the labour supply response at the time the first child becomes ineligible is larger for those with more remaining dependent children. This finding directly reveals that households are forward-looking, as labour supply responds contemporaneously to the increased salience of future benefit reductions.

Misattribution. Following Liebman and Zeckhauser (2004), a growing literature has documented evidence that people may misattribute the cause of changes to taxes or benefits. 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 change in marginal conditions (e.g. Feldman et al., 2016; Ito, 2012).

However, if households 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, (i) there would be no contemporaneous income effect and (ii) households should *reduce* labour supply when they become ineligible due to a perceived substitution effect.[23] The labour supply *increases* I have documented are therefore inconsistent with households having misattributed the benefit reduction to an increase in their marginal effective tax rate. Such an explanation also cannot rationalise the asymmetric responses between first and subsequent children or between large and small families.

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 increase; if they are prevented from doing so, there may be a discontinuous reduction in labour supply at the time of the increase.

Such an argument does not apply to a setting where people anticipate a future income *reduction* in unearned income, 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 4.

Discreteness in labour supply. The asymmetry in labour supply responses between the first and subsequent times a household's child ages out of eligibility for CTC could simply reflect discreteness in the labour supply decision: after mothers have increased labour supply the first time they lose eligibility, there may be little scope for further responses as other children become ineligible. Discreteness in labour supply alone cannot rationalise the pattern of labour supply responses, as a standard model would still predict no contemporaneous labour supply

[23] Indeed, this is Feldman et al. (2016)'s core argument and main finding.

response at the time of an anticipated benefit reduction. However, combining labour supply discreteness in a model with myopic or time inconsistent households could also rationalise (i) the existence of a labour supply response (driven by myopia or time inconsistency) and (ii) the asymmetry in response between first and subsequent children (driven by discreteness in labour supply preventing further labour supply increases at the time subsequent children become ineligible).

However, such an explanation would still be unable to rationalise the third empirical finding—that the labour supply response is larger among families with more remaining children, for which the increased salience of the eligibility rules generates a larger shock to lifetime income. This is because the frictions required to generate a labour supply response in a model with fully salient eligibility rules (such as myopia or time inconsistency) would not generate a relationship between the labour supply response and expected future reductions in benefit income. Therefore discreteness in labour supply (combined with some other friction) is also unable to rationalise all three features of the labour supply responses.

5 Robustness

In this section I set out evidence in support of my empirical strategy. In particular, I show that the labour supply of mothers in treatment group households follows the same trend in the months before losing eligibility as those in the control group (corroborating the visual evidence in Figure 2); and I also show that there are no differential changes in household composition, which could drive differential labour supply behaviour, between treatment and control groups at the time of the benefit reduction. These results reassure that my empirical strategy identifies a labour supply response to the benefit reduction. I then briefly discuss a number of additional robustness analyses contained in Online Appendix A.

5.1 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.

While it is not possible to test this assumption directly, I examine whether it holds in the period before becoming ineligible for CTC. Figure 2 provides visual evidence that this is the case. The labour supply in the lump-sum treatment group L evolves similarly to 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:

$$y_{it} = \alpha + \eta t + \sum_g \beta_g (T_i^g \times t) + \sum_g \gamma_g T_i^g + \zeta X_{it} + \theta_t + \epsilon_{it}. \tag{3}$$

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 A of Table 4 presents the estimates of β_g. All estimates 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. Overall, I find no evidence that my results are driven by differential trends in labour supply between the control and treatment groups.

Table 4: Robustness Tests

	All	First Child	Other Child	First Child, Large Family	First Child, Small Family	Other Child
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: 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)
Observations	43,209	43,209	43,209	43,209	43,209	43,209
Panel B: 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

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

5.2 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 B of Table 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.

5.3 Additional robustness tests

Online Appendix A contains a number of other robustness tests for the empirical results; I provide a brief overview of the results here. A key assumption in the analysis is that the time at which households become ineligible for CTC is exogenous because it is driven by the age of their children. However, for children aged between 16 and 20, whether their family remains eligible for CTC depends on whether the child remains in full-time non-advanced education.

In Appendix A.2, I show that the pattern and magnitude of estimated responses are robust to using the child's age as an instrument for eligibility. This reassures that the findings are not driven by endogeneity of the time households lose eligibility, driven by the child's choice about whether to remain in education.[24]

In Online Appendix A.3, I instrument for the household's treatment group assignment using a grouping estimator motivated by Blundell et al. (1998). The concern is that, if a household experiences a transitory 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. The pattern and magnitude of results is robust to this strategy, although the estimates are less precise.

Finally, I show in Online Appendix A.4 that the pattern of results is not affected by including single mothers in the sample, and in Online Appendix A.5 I show that there is no estimated labour supply response for men at the time of losing eligibility.

6 Model

The results in Section 4 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? And what are the welfare implications?

In this section, I outline a 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 then estimate key parameters within the model by the method of simulated moments using the reduced form results, which provide identifying variation, as targets. I use the estimated model to compute the welfare cost of being uninformed about benefit eligibility rules over the life-cycle, along with the implications for the government budget.

6.1 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.

[24] This is unsurprising because, in the period I study, there was near-universal uptake of post-compulsory education in the UK (which includes vocational qualifications). Almost all children therefore become ineligible at age 18. Additionally, children were eligible to claim means-tested payments, known as the Educational Maintenance Allowance, if they remained in education, providing support for children to continue studying following shocks to their parents' income.

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 risks. 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 that only the age of one child (the oldest) is required as a state variable, rather than the age of all children separately. Second, because the motivation for the model is labour supply responses among cohabiting women with older children, I do not consider marriage, divorce or endogenous fertility.

6.2 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.

6.2.1 Preferences

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

$$u(c_t, n_t) = \frac{\tilde{c}_t^{1-\gamma}}{1-\gamma} \exp\{U(n_t)\}, \tag{4}$$

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).[25] The parameter γ determines both risk aversion and the elasticity of intertemporal substitution and the function U, which determines how utility from consumption is affected by working, is given by

$$U(n_t) = \mathbb{1}[n_t = P]\alpha_P + \mathbb{1}[n_t = F]\alpha_F. \tag{5}$$

Since I estimate (1-γ) to be negative 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

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

$$V_t(X_t) = \max_{\{\tilde{c}_j, n_j\}_{j=t,\ldots,\bar{t}}} E_t\Big\{ \sum_{j=t}^{\bar{t}} \beta^{j-t} u(\tilde{c}_j, n_j) | X_t \Big\} \tag{6}$$

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 $\theta_t \in X_t$.

6.2.2 Budget constraint

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

$$a_{t+1} = (1+r)a_t + w_t n_t + w_t^m n_t^m + T(n_t, X_t) - c_t, \tag{7}$$

$$a_{t+1} \geq \underline{a}_{t+1}, \tag{8}$$

with the initial condition $a_0 = 0$ and terminal condition $a_{\bar{t}+1} = 0$. Households can save or 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.[26] 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

$$\ln w_t = \alpha_0 + \alpha_1 \ln(t-20) + v_t, \tag{9}$$

$$v_t = \rho v_{t-1} + \epsilon_t, \tag{10}$$

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

[26] 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.

$$\ln w_t^m = \alpha_0^m + \alpha_1^m \ln(t-20) + v_t^m, \tag{11}$$

$$v_t^m = \rho^m v_{t-1}^m + \epsilon_t^m, \tag{12}$$

$$\text{Prob}(n_t^m = F) = \text{Prob}(\omega_t^m > m_1 + m_2 \mathbb{1}[n_{t-1}^m = F]) \tag{13}$$

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. (2016), 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.[27] 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

$$k_t = \bar{k} - g(\bar{k}, t, t^k). \tag{14}$$

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) in Section 2.

6.2.3 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 salience parameter θ thinks that its number of dependent children in period t will be

$$\tilde{k}_t(\theta) = \bar{k} - \theta g(\bar{k}, t, t^k) \tag{15}$$

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.

[27] 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.

6.3 Discussion

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 salience *increases through experience*. In this second scenario, which is motivated by the reduced form results in Section 4, 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. That is, $\theta = 1$ if and only if one of the family's children has aged out of eligibility.

Full salience. Under full salience, a foreseeable reduction in benefits has no contemporaneous income effect because it is already included in the worker's marginal utility of lifetime wealth. There would therefore be no labour supply response to a purely lump-sum reduction in CTC.[28] Households adjust labour supply at the time of a foreseeable change in their benefits only if it affects the benefit withdrawal rate they face and hence their effective wage rate.

Salience increasing with experience. An increase in the salience of the eligibility rules, reflected by a change in θ, leads households to realise that lifetime unearned income will be lower than previously expected and therefore reduces the perceived marginal utility of lifetime wealth. Correspondingly, households would increase labour supply, consistent with the reduced form results in Section 4.

Salience increasing with experience would also generate other features of the reduced form results. First, having learned that benefit receipt depends child age, there should be no income effect at the time of any subsequent children becoming ineligible. This is because the household incorporates all future reductions in CTC into its marginal utility of lifetime wealth at the time the first child becomes ineligible—and so adjusts labour supply to enable it to smooth marginal utility at the time of all future benefit reductions. Second, the income effect on labour supply at the time of learning about the eligibility rules will be larger for mothers with more still-eligible children. These households 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. The model set out in this section is therefore consistent qualitatively with the reduced form results if the salience of the eligibility rules increases with experience.[29] I now estimate key parameters of the model and show that it is also capable of generating the empirical labour supply responses quantitatively.

6.4 Estimation

Target moments and identification. Within the structure of the model, the empirical results identify preference parameters (which determine the size of the labour supply response

[28] 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 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.

[29] And, as I discuss in Section 4.4, the reduced form results are inconsistent with a set of competing models.

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 initial average salience of the eligibility rules $\bar{\theta}$ 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. As the reduced form analysis shows there is only a labour supply adjustment when a household's first child ages out of eligibility (suggesting that the salience of the rules increases for all households following this initial experience), I impose in the model that $\theta = 1$ for all households after their eldest child becomes ineligible.

The identification argument is as follows. 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 because, given the female wage equation, preferences over full- and part-time work are key determinants of average earnings and the employment rate. Then, (iii) provides variation in the labour supply response to differently sized future income reductions, which are governed in the model by the elasticity of intertemporal substitution $1/\gamma$. Finally, given knowledge of the preference parameters which determine the size of a labour supply response to a household for which the eligibility rules were initially non-salient, (iv) provides information on the average number of households who were uninformed about the benefit reduction $(1 - \bar{\theta})$. This is simply 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 outside the model. I use a two-step procedure to estimate the model, first setting a number of parameters exogenously based on existing literature or characteristics of the lump-sum treatment group L which is the focus of this paper. 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 5.

The wage equation parameters draw on Blundell et al. (2016), who estimate male and female wage processes using the same data, and for a similar period, as my empirical analysis. I show the life-cycle profiles of average wages in Figure 3.

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 characteristics of the lump-sum treatment group L. 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 simulated sample in Table 6.

Table 5: 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

Notes: Table contains values of model parameters set outside the model. See text in Section 6.4 for details.

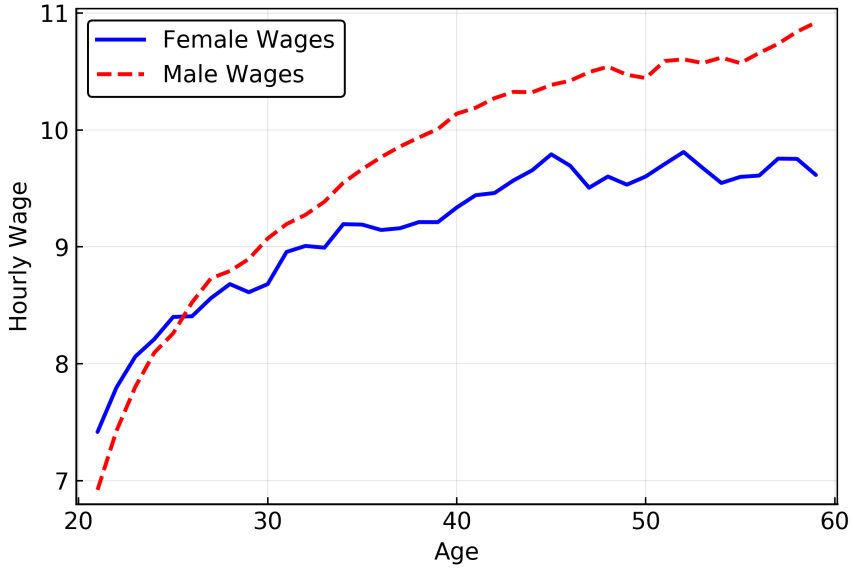

Figure 3: Male and Female Wage Processes

Notes: Chart compares average male and female wage realisations, in a simulated sample of 500 households, as a function of age.

These proportions are set so that the simulated sample matches the lump-sum treatment group L, and captures the empirical relationship between family size and age of children. Specifically, the proportions of households by number of children in the model are equal to the proportions in the sample. Then, within each family size, the ages at which the first child becomes ineligible are set equal to 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

Table 6: 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%**

Notes: Table shows the joint distribution of family size and household age when the first child ages out of eligibility. See text in Section 6.4 for details.

these groups to 40, 45 and 50 in the model.

Estimation procedure. I then estimate the four remaining parameters by the method of simulated moments: these are three preference parameters $\{\gamma, \alpha_P, \alpha_F\}$ and the average salience of the eligibility rules $\bar{\theta}$. I use the reduced form estimated labour supply responses as targets. To compute their equivalent in the model, I first 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 for each household, using the *same* household's behaviour in the full salience case as the counterfactual. For each household i, I calculate the change in female labour income at the time child k becomes ineligible under the full salience and non-salience scenarios: respectively, $\Delta y^S_{i,k}$ and $\Delta y^{NS}_{i,k}$. This captures that female labour supply may change predictably at the time a child becomes ineligible due to a direct effect of child age on female labour supply (arising through a change in the consumption equivalence scale in the model). Household i's labour supply response to the benefit reduction when child k becomes ineligible is given by the difference in differences:

$$\Delta y_{i,k} = \Delta y^{NS}_{i,k} - \Delta y^{S}_{i,k}. \tag{16}$$

I then average these household-specific treatment effects. I construct treatment groups analogously to the reduced form analysis, 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 3.1, 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

$$\Delta \bar{y}_g = \frac{1}{N_g} \sum_{\{i,k\} \in g} \Delta y_{i,k} \tag{17}$$

Table 7: 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

where N_g is the number of household-child pairs $\{i, k\}$ in group g.[30] This is the average labour supply response that would arise 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 4.

I also target 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,

$$\bar{y}_g = (1 - \bar{\theta})\bar{y}_g^{NS} + \bar{\theta}\bar{y}_g^{S}, \tag{18}$$

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

6.5 Parameter estimates and model fit

Table 7 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.

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 prevent 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.

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),

[30] 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 8: 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

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. (2016) for low-educated cohabiting 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 8. 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 4 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 reduced form analysis, designed to hold these predictable labour supply increases constant between groups. The difference in difference of the dashed and solid lines is the average labour supply response to increased salience of the eligibility rules.

Third, 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 6, these younger households have a higher composition of larger families. These families have a larger information shock in the learning scenario and so increase 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 future years in which the actual entitlement to CTC is below their initial (mistaken) expectation.

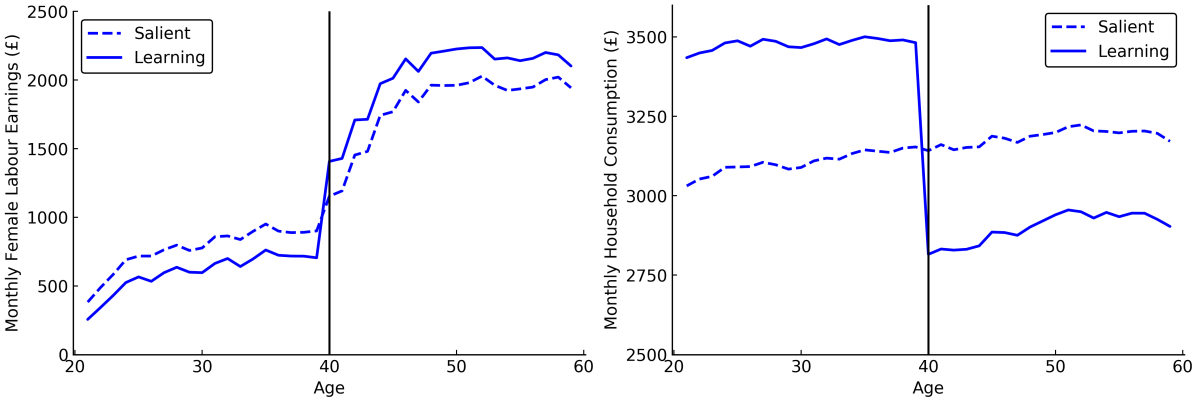

(a) Female Labour Income (lose CTC age 40)

(b) Household Consumption (lose CTC age 40)

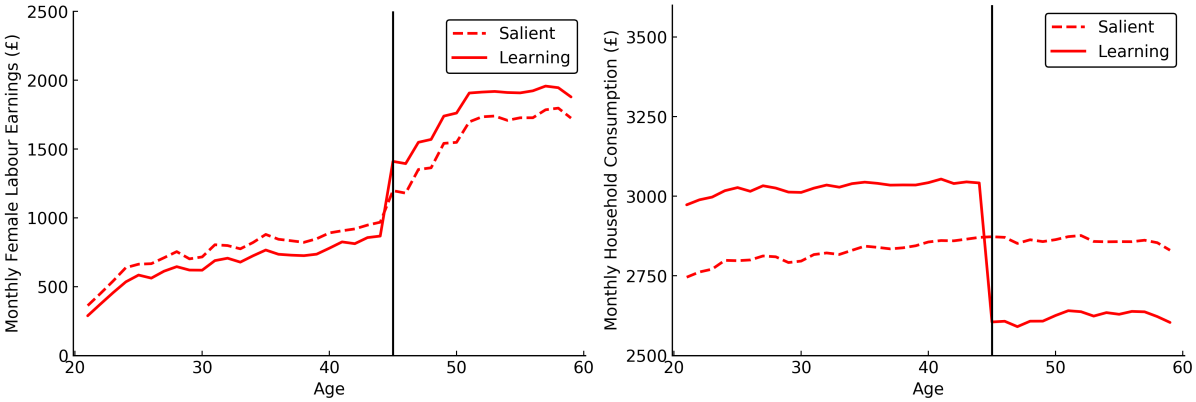

(c) Female Labour Income (lose CTC age 45)

(d) Household Consumption (lose CTC age 45)

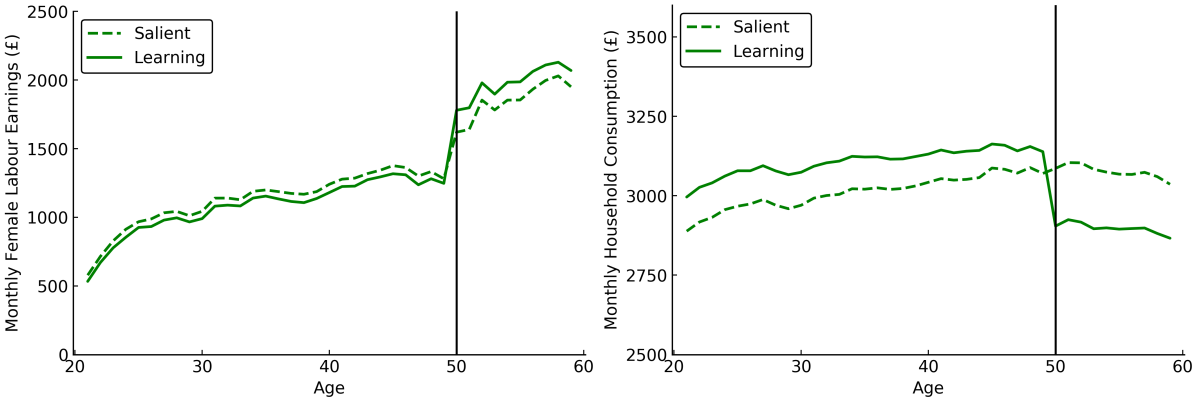

(e) Female Labour Income (lose CTC age 50)

(f) Household Consumption (lose CTC age 50)

Figure 4: Lifecycle Profiles of Monthly Labour Income and Consumption

Notes: Charts compare profiles of monthly labour income (left column) and household consumption (right column) over the life-cycle as implied by the model. The rows differ by the age of the household when the first child becomes ineligible. Each panel shows the difference between profiles if the household is fully aware of the eligibility rules and if they only learn about them through experience.

6.6 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 initial non-salience (EV_0^{NS}), allowing lifecycle consumption, savings and labour

supply choices to adjust:

$$EV_0^S(\Delta A_0) = EV_0^{NS}(0) \tag{19}$$

I show the compensating variation in row (1) of Table 9, separately by family size and as a weighted average across all families. Larger families experience 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 fully aware of the benefit eligibility rules (rather than unaware) across their entire lifecycle. 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.[31]

To provide a measure of the size of the welfare cost, I express average welfare cost as a proportion of total lifetime CTC receipts. 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 are substantial. There are two reasons for this: first, the financial impact of losing eligibility for CTC are large. But second, and importantly, the age-related eligibility rules provide dynamic incentives for labour supply (and consumption) over the entire lifecycle. Failing to anticipate the benefit reduction therefore leads to optimisation errors which compound over time, necessitating especially costly adjustment when the benefits fall.

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 4, the life-cycle profile of labour supply depends strongly on households' awareness in the eligibility rules, and 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.

[31] 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 9: 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

Notes: Table contains the estimated average welfare costs of the non-salient eligibility rules. See the text in Section 6.6 for details.

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

6.7 Summary and discussion

I now briefly discuss the main conclusions from the structural analysis and emphasise the key results. First, using the empirical results from Section 4 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 salience regimes—closely replicates the empirical results from Section 4. This suggests that my proposed mechanism of initial non-salience of the eligibility rules (which increases only after a household's first child becomes ineligible) 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, as they may be highly non-salient (even when the financial stakes are high), generating substantial welfare costs. This points to major source of inefficiency in the welfare system which has not previously been documented.

7 Conclusion

This paper 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 anticipate future benefit reductions. Using a structural life-cycle model of labour supply, I estimate that the proportion of claimants initially unaware of the benefit eligibility rules is high. And I find that the welfare costs to claimants from being uninformed about the eligibility rules are substantial. Taken together, the results highlight an important source of inefficiency in the welfare system which has not previously been documented.

The results underline the potential for incentives to be non-salient, even in real-world systems and even when the financial stakes are high. 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. For example, individuals may be unaware of features of their pension schemes until they retire, or of insurance coverage until they make a claim. Future research could usefully investigate whether this is the case, as the costs of making errors in these environments are potentially substantial.

References

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.

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. (2016). 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*, 66(4):827–861.

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

Brewer, M. (2003). The new tax credits. IFS Briefing Note 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.

Chetty, R. and Saez, E. (2013). Teaching the Tax Code: Earnings Responses to an Experiment with EITC Recipients. *American Economic Journal: Applied Economics*, 5(1):1–31.

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

Dalton, C. M., Gowrisankaran, G., and Town, R. J. (2020). Salience, Myopia, and Complex Dynamic Incentives: Evidence from Medicare Part D. *The Review of Economic Studies*, 87(2):822–869.

Department for Work and Pensions (2015). 2010 to 2015 government policy: welfare reform. Available at: https://www.gov.uk/government/publications/2010-to-2015-government-policy-welfare-reform/2010-to-2015-government-policy-welfare-reform.

Fang, H. and Wang, Y. (2015). Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. *International Economic Review*, 56(2):565–596.

Farhi, E. and Gabaix, X. (2020). Optimal Taxation with Behavioral Agents. *American Economic Review*, 110(1):298–336.

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.

Goldin, J. (2015). Optimal tax salience. *Journal of Public Economics*, 131:115–123.

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

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.

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

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

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. Congressional Budget Office Working Paper 2012-12.

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.

Paserman, M. D. (2008). Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation*. *The Economic Journal*, 118(531):1418–1452.

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

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.

Taubinsky, D. and Rees-Jones, A. (2017). Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment. *The Review of Economic Studies*, 85(4):2462–2496.

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.

Online Appendix

A Additional Results

A.1 Extensive margin responses

The main results in Section 4 focus on mothers' monthly labour income as the measure of their labour supply. This section presents estimates of equivalent specifications which instead adopt a binary measure of labour supply (either working or not) as the dependent variable. The overall pattern of the point estimates is similar, although the standard errors on the estimates are larger.

Column (1) shows the estimated treatment effect for all mothers. The extensive margin labour supply response is small compared with the labour income response documented in Section 4.1, with a statistically insignificant increase of 2.1 percentage points (s.e. 2.0). Taken together with the results in Section 4.1, these results point to a substantial intensive margin response, with mothers who already work increasing labour income in response to losing CTC.

Table 10: 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)
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

Notes: Table shows 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. Column (1) shows the treatment effect estimated at the time any child ages out of eligibility; Columns (2) and (3) distinguish between first and subsequent children; Columns (4) to (6) distinguish first children from large vs. small families. Panel B expresses the labour supply treatment effects as elasticities. Standard errors, clustered at the level of the mother, in parentheses.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Columns (2) and (3) demonstrate that, as with the responses for labour income, all of the extensive margin response is driven by families losing eligibility for the first time. And columns (4) to (6) demonstrate that the point estimates for the labour supply response at the extensive margin for households whose first child becomes ineligible is also larger for households with more remaining children (although the estimate for large families is not statistically significant). Overall, the pattern of point estimates at the extensive margin is consistent with the estimates for monthly labour income presented in the main text.

A.2 Exogeneity of losing eligibility

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 (2) 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,[32] 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 ineligibility, I repeat each analysis using the child's age as an instrument for their eligibility. This strategy is designed to use only the exogenous 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

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

and use the predicted values $\hat{A}_{it}^{IV}$ in place of A_{it} in each of the three analyses described in Section 4. I present the estimates for each analysis in Table 11.[33] While these IV estimates of the labour supply effects are less precise than the baseline estimates in Section 4, 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 prob-

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

[33] 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 4. 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 11: Instrumenting for A_{it}

	All	First Child	Other Child	First Child, Large Family	First Child, Small Family	Other Child
	(1)	(2)	(3)	(4)	(5)	(6)
Labour Income (£/mo.)	114.6 (157.2)	181.4 (157.3)	-24.75 (167.7)	239.7 (283.9)	179.8 (158.6)	-17.40 (165.8)
Probability of working	0.0166 (0.0516)	0.0404 (0.0534)	-0.0332 (0.0627)	0.142 (0.144)	0.0339 (0.0526)	-0.0326 (0.0579)
Observations	62,754	62,754	62,754	62,754	62,754	62,754

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

ability 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 3.

Next, columns (2) and (3) in Panel A of Table 11 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 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 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 £239.7 a month, and increase the probability of working of 14.2 percentage points. The IV estimated labour supply responses for mothers with fewer remaining children in column (5) are smaller at £179.8 and 3.4 percentage points. 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 4 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.

A.3 Treatment group assignment

As I describe in Section 3.1, 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.As I describe in Section 3.1, 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.

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 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

$$W_{it} = \alpha + G_i'\phi_1 + (G_i \times A_{it})'\phi_2 + X_{it}'\zeta + v_{it} \tag{21}$$

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 Table 12. As with the IV analysis for eligibility status, the magnitude and pattern of the point estimates

Table 12: Instrumenting for T_i

	All	First Child	Other Child	First Child, Large Family	First Child, Small Family	Other Child
	(1)	(2)	(3)	(4)	(5)	(6)
Labour Income (£/mo.)	125.2 (119.1)	151.7 (130.8)	71.82 (156.5)	338.3 (256.7)	140.4 (127.3)	121.3 (149.1)
Probability of working	0.0837* (0.0455)	0.119** (0.0572)	0.0377 (0.0751)	0.164 (0.106)	0.0820 (0.0518)	0.0806 (0.0654)
Observations	72,531	72,531	72,531	72,531	72,531	72,531

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

are very similar to the baseline estimates across all three analyses, albeit estimated less precisely as reflected by 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 4 are not driven by endogeneity in treatment group assignment. The pattern of results in Table 12 is the same as the main results in Section 4, and remain consistent with CTC claimants learning about initially non-salient eligibility rules through experience.

A.4 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 cohabiting mothers are also present for this group.

In Table 13 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 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.

Table 13: Estimates including single mothers in the sample

	All	First Child	Other Child	First Child Lage Family	First Child Small Family	Other Child
	(1)	(2)	(3)	(4)	(5)	(6)
Labour income (£/ month)	163.4*** (59.60)	190.6*** (61.37)	95.4 (63.94)	288.47*** (93.30)	184.90*** (61.74)	96.81 (61.16)
Probability of working	0.033* (0.018)	0.043** (0.019)	0.007 (0.025)	0.042 (0.057)	0.044** (0.019)	0.006 (0.023)
CTC amount (£/ month)	-151.79*** (1.43)	-150.47*** (1.44)	-155.04*** (1.56)	-155.11*** (2.54)	-150.96*** (1.44)	-153.38*** (1.60)
Taper rate	-0.000103 (0.000436)	-0.00008 (0.00043)	-0.00024 (0.00071)	0.00087 (0.00105)	-0.00048 (0.00048)	0.00058 (0.0008)
Observations	109,042	109,042	109,042	109,042	109,042	109,042

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

A.5 Labour supply responses among men

I focus on any response to female labour supply at the time of the foreseeable reduction in CTC as a means to infer the exten to to which the reduction was anticipated. This is because female labour supply is typically found to be more responsive than male labour supply (e.g. Blundell and MaCurdy, 1999; Meghir and Phillips, 2010). In Table 14, I demonstrate that I indeed find no statistically significant labour supply response among males at the time of the reduction in CTC eligibility.

Table 14: Labour supply response for men

	All
Treatment Effects (β_L)	
Labour income (£/ month)	94.1 (65.38)
Probability of working	0.023 (0.025)
CTC amount (£/ month)	-141.7*** (3.06)
Taper rate	0.00104 (0.00073)
Observations	38,542

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

B Data Appendix

In this appendix, I provide further details on the way I constructed the data used in the paper. 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).

B.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).[34]

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

[34]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.

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

B.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.[35] 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

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 $\bar{Y}_1$.[36]

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 (2) 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 3.1 I describe the three main groups used in the analysis: L, M and C. I assigned mothers to treatment groups using the following procedure:[37]

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 $\tau(k_t - 1, Y_t)$;

[35] I treat a part-time job as 20 hours a week, and full-time as 40 hours.
[36] Over the period 2003 to 2013, $\bar{Y}_0$ increased from £5,060 to £6,420.
[37] 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.

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