

Shaw, Jonathan

Working Paper

The redistribution and insurance value of welfare reform

IFS Working Papers, No. W14/21

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Shaw, Jonathan (2014) : The redistribution and insurance value of welfare reform, IFS Working Papers, No. W14/21, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/wp.ifs.2014.1421>

This Version is available at:

<https://hdl.handle.net/10419/101339>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The redistribution and insurance value of welfare reform

IFS Working Paper W14/21

Jonathan Shaw

The author gratefully acknowledges a grant from the Nuffield Foundation (OPD/40976) and co-funding from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy at IFS (RES-544-28-5001).

The Institute for Fiscal Studies (IFS) is an independent research institute whose remit is to carry out rigorous economic research into public policy and to disseminate the findings of this research. IFS receives generous support from the Economic and Social Research Council, in particular via the ESRC Centre for the Microeconomic Analysis of Public Policy (CPP). The content of our working papers is the work of their authors and does not necessarily represent the views of IFS research staff or affiliates.

The redistribution and insurance value of welfare reform

Jonathan Shaw*

August 12, 2014

Abstract

Relatively little is known about the roles that taxes and transfers play in redistributing resources and providing insurance across individuals and across the lifecycle. We embed these alternative roles in a lifecycle model, allowing us to demonstrate what the tax and transfer system achieves from a lifecycle perspective and why it is valuable. We undertake a five-way decomposition of net transfers into a giveaway term and terms corresponding to between- and within-individual redistribution and between- and within-individual insurance. These components are distinguished from perspective of the start of working life, and we consider both the magnitude of net transfers involved and the associated welfare values. Our focus is on females and we also highlight how behavioural responses affect the results. Analysis is conducted for the 2015 UK tax and transfer system relative to a flat-rate baseline, showing what value is provided by the complex tax and welfare entitlement rules in a modern economy. We also consider what is achieved by two important UK benefit reforms—the working families’ tax credit (WFTC) reform of 1999 and the universal credit (UC) reform that began in 2013. Our main conclusions are that insurance against wage and family composition shocks is substantial and highly valued by individuals. Within-individual redistribution (i.e. across periods of life) is generally of little value even in the presence of strict borrowing constraints. Behavioural responses tend to increase the size of reform giveaways at the expense of the other components.

1 Introduction

It is common to assess tax and transfer systems from a snapshot perspective using cross-sectional data. For example, progressivity is often measured using short-run Gini-based measures of redistribution or indicators of how average

*IFS and UCL. I gratefully acknowledge a grant from the Nuffield Foundation (OPD/40976) and co-funding from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy at IFS (RES-544-28-5001). I have benefitted significantly from discussions with Richard Blundell, Eric French and Magne Mogstad. All remaining errors are my own.

annual tax rates change across the income distribution; see e.g. Office (2012) and Paturot et al. (2013). Studies comparing across countries often use these types of cross-sectional measure to distinguish between different types of tax and benefit system, contrasting the progressive “Nordic model” countries with the less progressive “Anglo-Saxon model” countries such as the UK and US.

But something that looks like redistribution at a point in time may appear very different from the perspective of the lifecycle. For example, consider parental leave following the birth of a child. This can largely be thought of as a policy of enforced saving, whereby the individuals concerned pay for the parental leave benefits they receive through the taxes they pay while at work. A static perspective also ignores the insurance value provided by taxes and benefits. Individuals face considerable uncertainty over future circumstances, and one important role of the tax and benefit system is to insure against that uncertainty.

We know very little about how much—and what sorts of—redistribution and insurance tax and transfer systems provide from a lifecycle perspective. Indeed, it is even possible that Anglo-Saxon countries such as the US turn out to have more progressive systems than some nordic model countries. We also know little about how valuable the redistribution and insurance components are to individuals.

In this paper, we take an important step towards a better understanding of the roles played by taxes and transfers across the lifecycle. We exploit a lifecycle model developed in Blundell et al. (2013) and build on the decomposition methodology set out in Hoynes and Luttmer (2011) to decompose net transfers into five components: an overall giveaway term, redistribution between individuals, redistribution within individuals (i.e. across periods of life), insurance between individuals and insurance within individuals. On the basis of this decomposition we are able to demonstrate what the tax and transfer system achieves from a lifecycle perspective and why it is valuable. Components of the decomposition are measured from perspective of the start of working life, and our focus is on females. We also highlight how behavioural responses affect the results.

We conduct analysis for the 2015 UK tax and transfer system relative to a flat-rate baseline, showing what is achieved by the complex tax and welfare entitlement rules in a modern economy. We also consider the impact of two important UK benefit reforms—the working families’ tax credit (WFTC) reform of 1999 and the universal credit (UC) reform that began in 2013.

Our main conclusions are that insurance against wage and family composition shocks is substantial and highly valued by individuals. Within-individual redistribution (i.e. across periods of life) is generally of little value even in the presence of strict borrowing constraints. Behavioural responses tend to increase the size of reform giveaways at the expense of the other components.

2 Literature

This paper is related to three existing strands of literature. The first strand is the set of papers evaluating the impact of work-contingent transfer programmes such as the WFTC reform in the UK that we will consider and Earned Income Tax Credit (EITC) reforms in the US. Blundell et al. (2005) is a prominent example addressing the WFTC reform. In common with many of the other papers in this area, the outcome of interest is employment and estimated impacts are derived using a difference-in-differences approach. They find that the reform increased the employment rate of lone parents by 3.6 percentage points, with little effect on couples with children. Papers doing the same for the EITC reforms in the US include Eissa and Liebman (1996), Meyer and Rosenbaum (2001) and Eissa and Hoynes (2006). Francesconi and Van der Klaauw (2007) and Francesconi et al. (2009) use the same difference-in-differences approach to evaluate the WFTC reform, but consider a wider range of outcomes, including childcare usage, marriage, divorce and fertility.

The disadvantage of a difference-in-differences approach is that it doesn't provide the necessary framework to evaluate welfare changes; for this, a structural approach is needed. Brewer et al. (2006) take such an approach to evaluating the WFTC reform, focusing on labour market impacts and the utility cost of programme participation. They find the reform increased labour supply of lone mothers by 5.1 percentage points with relatively small impacts for other groups. For lone mothers, they find that part of this impact is due to a fall in the utility cost of participation. As far as we are aware, Blundell et al. (2013) is the only paper that has tried to estimate the impact of WFTC in a dynamic framework with uncertainty that allows redistribution and insurance to be distinguished. We build on that model to undertake a formal decomposition of the WFTC reform into four components: redistribution and insurance each between and within individuals.

The second strand of literature addresses lifecycle inequality and redistribution, investigating how the impression of inequality and redistribution done by taxes and benefits changes when one takes a lifecycle rather than an annual perspective. For example, Bovenberg et al. (2008) calculates what fraction of benefits received by individuals is self-financed (by taxes they pay in the same or other periods of life). They decompose total benefit expenditure into three components: within-year within-person redistribution, across-year within-person redistribution and across-person redistribution. For Denmark, they find that three-quarters of a selected set of public transfers are self-financed. This exercise has been repeated for other countries by, among others, Falkingham and Harding (1996) and O'Donoghue (2002).

Another set of papers in this strand has focused on the progressivity of lifecycle redistribution by comparing gross and net income inequality. For example, Björklund and Palme (1997) apply a measure of inequality that can be decomposed into "within-group" and "between-group" components to a dataset of individuals across all periods of their adult life. The within-group component corresponds to within-person inequality and the between-group component

corresponds to across-person inequality. By calculating this for net and gross income and taking the difference between corresponding terms, they arrive at measures of within- and across-person redistribution. Applying this to Swedish data, they find that taxes and benefits reduce overall inequality by between 20 and 35 per cent. Across-person inequality falls more than within-person inequality, but the split is sensitive to the treatment of children and the degree of inequality aversion. They also show that income taxes account for most of the fall in across-person inequality, while welfare benefits are largely responsible for the reduction in within-person inequality – presumably because benefits tend to be focused more towards times of temporary hardship. Bartels (2011) applies the same technique to Germany.

All these papers in the second strand of literature take an ex post perspective, counting all taxes and benefits as redistribution. We build on this work by taking an ex ante perspective, allowing a distinction between redistribution and insurance to be made.

The third strand of literature investigates the impact and welfare value of different social insurance programmes. These include Card et al (2009) on health, Low and Pistaferri (2010) on disability, Gruber (2000) and Chetty (2008) on unemployment, Low et al. (2010) and Hoynes and Luttmer (2011) on US state taxes and transfers.

These last two papers are closest in spirit to the present one. Low et al. (2010) study the welfare cost of risk and estimate the value of insurance provided by different types of benefits. Their results are derived from a lifecycle model of consumption and labour supply that distinguishes between employment risk (job arrival and job destruction) and productivity risk (permanent shocks to wages). The welfare cost of productivity risk is shown to be much higher than employment risk because the former has a more persistent effect on income. Consequently, means-tested benefits (which partially insure against productivity risk) have a much greater welfare value than does unemployment benefit (which partially insures against employment risk).

Hoynes and Luttmer (2011) decompose the value of US state taxes and transfers into redistributive and insurance components. They find that the value of redistribution provided by state tax and transfer systems falls sharply with income (unsurprisingly) but the value of insurance is positive across the income distribution and actually increases with income. Overall, three quarters of the population are better off for the existence of the state tax and transfer system. The decomposition we use develops the one set out in that paper, showing how it can be implemented in a model that includes behavioural responses and allowing all four components to be calculated rather than the two they are restricted to due to modelling limitations.

3 Measurement

We now describe how we will measure redistribution and insurance. Let $\Gamma(\Psi_a, y_a)$ be the annual net tax function (taxes less benefits). The arguments of this

function are the individual's information set, Ψ_a (e.g. age, education, family composition, etc) and her choices $y_a = \{c_a, l_a\}$ over consumption and labour supply, where a denotes the current age. We suppress the individual index i for the time being.

All decompositions are taken from the perspective of the start of adult life (after the end of full-time education), meaning that this is the age at which the distinction between redistribution and insurance is made. This seems a natural perspective against which to analyse the tax and transfer system. Define the value of expected lifetime utility at the start of adult life, age \underline{a} , as $U(y(\Psi, \Gamma) | \Psi_{\underline{a}}, \Gamma)$. In this, y is a vector of choice pairs, one for each period of life, shown as a function of the vector of information sets across life, Ψ . Now consider a reform to the net tax function from a baseline system Γ^0 to a reform system Γ^1 . Viewing things from the start of adult life, we can estimate how much the individual values the reform across the lifecycle by calculating proportional increment to consumption, π_{EV} , she would be willing to accept in place of the reform, holding labour supply fixed:

$$U(c(\Psi, \Gamma^0) \times (1 + \pi_{EV}), l(\Psi, \Gamma^0) | \Psi_{\underline{a}}, \Gamma^0) = U(y(\Psi, \Gamma^1) | \Psi_{\underline{a}}, \Gamma^1) \quad (1)$$

This is a form of consumption equivalent variation. It is similar to the measure Low et al. (2010) use to estimate the welfare cost of risk, except that they increment reform consumption rather than baseline consumption (making it a compensating variation rather than an equivalent variation). With the CRRA form of utility used below, it is possible to solve explicitly for the value π_{EV} :

$$\pi_{EV} = \left[\frac{U(y(\Psi, \Gamma^1) | \Psi_{\underline{a}}, \Gamma^1)}{U(y(\Psi, \Gamma^0) | \Psi_{\underline{a}}, \Gamma^0)} \right]^{\frac{1}{1-\gamma}} - 1$$

Calculating averages of this across heterogeneous individuals requires strong assumptions—see Fleurbaey (2009).

The change in net tax function from Γ^0 to Γ^1 implies a change in the amount of redistribution and insurance provided by taxes and benefits, the values of which we can separate. To do this we need to make precise our definitions of redistribution and insurance. We use *redistribution* to mean the expected value of net taxes conditional on current information (or, more informally, net taxes that are predictable on the basis of current information). We use *insurance* to mean the difference between actual and expected net taxes conditional on current information (or net taxes that are not predictable on the basis of current information).

This distinction allows us to write variation of net taxes around its expected value as the sum of redistribution and insurance components:

$$\Gamma - E[\Gamma] = \underbrace{E[\Gamma | \Psi_{\underline{a}}, a] - E[\Gamma]}_{\text{Redistribution}} + \underbrace{\Gamma - E[\Gamma | \Psi_{\underline{a}}, a]}_{\text{Insurance}}$$

where for simplicity we suppress the arguments of the net tax function, Γ . Each of these components can be further subdivided depending on whether the

variation is between or within individuals:

$$\begin{aligned}
\Gamma - E[\Gamma] &= \underbrace{E[\Gamma | \Psi_a] - E[\Gamma]}_{\text{Between individual}} + \underbrace{E[\Gamma | \Psi_a, a] - E[\Gamma | \Psi_a]}_{\text{Within individual}} \\
&\quad \underbrace{\hspace{10em}}_{\text{Redistribution}} \\
&+ \underbrace{E\{\Gamma - E[\Gamma | \Psi_a, a] | \Psi\}}_{\text{Between individual}} + \underbrace{\Gamma - E[\Gamma | \Psi_a, a] - E\{\Gamma - E[\Gamma | \Psi_a, a] | \Psi\}}_{\text{Within individual}} \\
&\quad \underbrace{\hspace{10em}}_{\text{Insurance}}
\end{aligned}$$

Now, based on this decomposition, let us define the following terms:

$$\begin{aligned}
\gamma^{RB}(\Psi_a) &= E[\Gamma | \Psi_a] - E[\Gamma] \\
\gamma^{RW}(\Psi_a, a) &= E[\Gamma | \Psi_a, a] - E[\Gamma | \Psi_a] \\
\gamma^{IB}(\Psi) &= E\{\Gamma - E[\Gamma | \Psi_a, a] | \Psi\} \\
\gamma^{IW}(\Psi, a) &= \Gamma - E[\Gamma | \Psi_a, a] - E\{\Gamma - E[\Gamma | \Psi_a, a] | \Psi\}
\end{aligned}$$

where the R superscript denotes redistribution, I insurance, B between and W within. Using these definitions, we can now define a sequence of intermediate tax and benefit systems between our baseline and reform tax and benefit systems, Γ^0 and Γ^1 , each of which does one more component of the reform than the previous one.

Γ^0	Baseline system
$\Gamma^N = \Gamma^0 + E[\Gamma^1] - E[\Gamma^0]$	Add reform giveaway
$\Gamma^{RB} = \Gamma^N + \gamma^{RB1}(\Psi_a) - \gamma^{RB0}(\Psi_a)$	Add change in between redistribution
$\Gamma^R = \Gamma^{RB} + \gamma^{RW1}(\Psi_a, a) - \gamma^{RW0}(\Psi_a, a)$	Add change in within redistribution
$\Gamma^{RIB} = \Gamma^R + \gamma^{IB1}(\Psi) - \gamma^{IB0}(\Psi)$	Add change in between insurance
$\Gamma^1 = \Gamma^{RIB} + \gamma^{IW1}(\Psi, a) - \gamma^{IW0}(\Psi, a)$	Add change in within insurance

This gives a sequence of intermediate tax and benefit systems that will allow us to separate the different redistribution and insurance components using equivalent variations base on equation (1). Specifically, we can write:

$$\begin{aligned}
U(c(\Psi, \Gamma^0) \times (1 + \pi_{EV}^N), l(\Psi, \Gamma^0) | \Psi_a, \Gamma^0) &= U(y(\Psi, \Gamma^N) | \Psi_a, \Gamma^N) \\
U(c(\Psi, \Gamma^N) \times (1 + \pi_{EV}^{RB}), l(\Psi, \Gamma^N) | \Psi_a, \Gamma^N) &= U(y(\Psi, \Gamma^{RB}) | \Psi_a, \Gamma^{RB}) \\
U(c(\Psi, \Gamma^{RB}) \times (1 + \pi_{EV}^{RW}), l(\Psi, \Gamma^{RB}) | \Psi_a, \Gamma^{RB}) &= U(y(\Psi, \Gamma^R) | \Psi_a, \Gamma^R) \\
U(c(\Psi, \Gamma^R) \times (1 + \pi_{EV}^{IB}), l(\Psi, \Gamma^R) | \Psi_a, \Gamma^R) &= U(y(\Psi, \Gamma^{RIB}) | \Psi_a, \Gamma^{RIB}) \\
U(c(\Psi, \Gamma^{RIB}) \times (1 + \pi_{EV}^{IW}), l(\Psi, \Gamma^{RIB}) | \Psi_a, \Gamma^{RIB}) &= U(y(\Psi, \Gamma^1) | \Psi_a, \Gamma^1)
\end{aligned}$$

The order in which components are evaluated is not unique, so it is possible that the valuations we recover will depend on the particular ordering we have chosen. Nevertheless, we imagine that the difference made will be relatively minor, particularly for the reforms we consider.

So far we have abstracted from issues relating to the family. Because families change over time, the only coherent way to analyse lifecycle issues is to follow individuals. But taxes and benefits in many countries depend to some degree on family income and circumstances. As a result, we need some way of apportioning taxes and benefits to individuals in couples. We take the simple approach of assuming equal sharing between members of a couple, i.e. all taxes paid and benefits received are split 50-50 between partners. As described below, we focus on females in our analysis. For consistency with the monetary decompositions, utility valuations hold partner net income fixed at the baseline level.

4 Model

So far we have set out the framework we will use to separate redistribution and insurance. We now describe the model we will use to implement this framework.

4.1 Model description

To estimate redistribution and insurance provided by the tax and benefit system, we use simulated data derived from a lifecycle model of female education, labour supply and savings decisions. A model of this type is needed for us to be able to value of each of the redistribution and insurance components. The model we use was developed in previous work (Blundell et al. (2013)). Here we provide an overview of the model, but see that paper for a comprehensive description.

The model is a lifecycle consumption and labour supply model following the tradition of Eckstein and Wolpin (1989) and Eckstein and Wolpin (1999), Keane and Wolpin (1997) and Adda et al. (2007), but with new features added to make it suitable for measuring redistribution and insurance provided by taxes and benefits. The key additional features are: accumulation of experience, evolving family composition and a detailed characterisation of UK taxes and benefits. As far as we are aware, this is the first capable of analysing the UK tax and benefit system from a dynamic perspective, allowing for labour supply responses, human capital formation and saving. We focus in the model on females since they are particularly vulnerable to family-related shocks such as the risk of lone-motherhood and also the need to take time out of the labour-market or work part time in order to look after children. Moreover, previous work has shown that females are often more responsive to tax and benefit changes. But the model allows the role of family composition to be considered: males are included as partners and we also model the arrival and departure of children.

Females are modelled on an annual basis between the ages of 17 and 69. Life is divided into three stages: education (decisions at age 17), working life (from labour market entry until age 59) and retirement (ages 60-69). First we describe working life. At each age a , a female i chooses consumption c_{ia} and labour supply l_{ia} to maximise her expected discounted lifetime utility given her current state defined by $\Psi_{ia} = (X_{ia}, \Omega_{ia}, \Pi)$. X_{ia} is observed idiosyncratic information, Ω_{ia} is unobserved idiosyncratic information and Π contains prices

and parameters for the tax and benefit system. The intertemporal problem can be written as:

$$V_{ia}(X_{ia}, \Omega_{ia}, \Pi) = \max_{\{c, l\}} E \left\{ \sum_{\alpha=a}^{\bar{a}} \beta^{\alpha-a} u(c_{i\alpha}, l_{i\alpha}; X_{i\alpha}, \Omega_{i\alpha}) \mid X_{i\alpha}, \Omega_{i\alpha}, \Pi \right\}$$

where the maximisation is subject to the standard intertemporal budget constraint:

$$k_{i,a+1} = Rk_{ia} + w_{ia}l_{ia} - \Gamma(\Psi_{ia}, y_{ia}) - c_{ia}$$

as well as constraints describing how other state variables evolve, discussed shortly. Period utility is assumed to take the following CRRA form:

$$u(c_{ia}, l_{ia}; X_{ia}, \Omega_{ia}) = \frac{(c_{ia}/n_{ia})^{1-\gamma}}{1-\gamma} \exp \{ \tilde{u}(s_{ia}, l_{ia}, d_{ia}^m, l_{ia}^m, d_{ia}^k, a_{ia}^k) + \theta_i(l_{ia}) \}$$

where n is an equivalence scale (adjusting for family size and composition), $\theta_i(\cdot)$ is the woman's unobserved permanent preference for each labour supply level, and $\tilde{u}(\cdot)$ is a utility shifter that depends on family characteristics. These family characteristics are education, s , labour supply, l , presence and labour supply of partner, d^m and l^m , and presence and age of youngest child, d^k and a^k . The female wage process is a fairly standard setup that depends on education, experience and a stochastic AR(1) productivity process intended to capture other factors affecting wages e.g. health shocks. We write this as:

$$\begin{aligned} \ln w_{ia} &= \ln W_s + \alpha_s \ln(e_{ia} + 1) + v_{sia} + \xi_{sia} \\ v_{sia} &= \rho_s v_{sia-1} + \epsilon_{sia} \end{aligned}$$

where w is the wage, W is the market wage rate corresponding to zero experience, e is experience and ξ is a transitory wage shock interpreted as measurement error that does not affect choices. v is productivity that follows an AR(1) process with iid innovations ϵ . Women accumulate one unit of experience for each year of full-time work and a fraction of a unit for each year of part-time work. Experience gradually depreciates during life regardless of work decisions. This is summarised by:

$$e_{i,a+1} = e_{ia}(1 - \lambda_{D_s}) + 1(l_{ia} = 40) + \lambda_{P_s} 1(l_{ia} = 20)$$

where λ_{D_s} and λ_{P_s} are rates of experience depreciation and accumulation, and $1(\cdot)$ is the indicator function.

Family composition is assumed to vary exogenously over time, with transition probabilities estimated to match those observed in data. Treating family composition as stochastic is not that common in this sort of model, but it may not be unreasonable. No one can predict accurately their future family circumstances, particularly when young. And data from the UK Millennium Cohort Study suggests that over 40 per cent of children born are the result of unplanned

pregnancies. We provide an informal account of the family formation process because the maths quickly becomes tedious; see Blundell et al. (2013) for the details. First partners: the probability that a single female finds a partner of a given education level depends on her age and education. Partners leave with a probability that depends on the partner’s education, the female’s age and the presence of a child. Partners are fully characterised by: their education, their employment status and their wage. They work either full time (40 hours per week) or not at all and wages follow a similar process as that for females, except that for simplicity age replaces experience (and parameters values are allowed to differ). Female decisions are made conditional on those of any partner she has.

Turning to children, the probability that a child is born depends on age and education of the female, presence of partner and of older children. Children depart for sure after age 18. To prevent the state space becoming too large, we keep track of the age of youngest child but not number of children in the family. This is because previous work has shown that age of youngest child is what matters for work decisions. We model a fraction of families as having access to free childcare (e.g. from grandparents), with the remainder needing to pay for childcare if no one stays at home to look after the children. The need to pay for childcare is drawn randomly at the start of life, and the fraction that pays is estimated from the data. The required number of hours of childcare varies by age of child: children aged five or under need childcare for every hour where all adults in the family are out working; children aged 6-10 only need childcare outside school hours, and children aged 11 or older are assumed not to need childcare. This is broadly in line with what we find in the data. The hourly price for childcare is a uniform average calculated from the data.

Finally, we assume that a fixed share of families (calculated separately for each education level) rent their home throughout life, with rent set equal to the median in the data (which doesn’t vary much by family type).¹ Given family circumstances and gross earnings of all adults in the family, we calculate net family income using FORTAX, as described above. We equalise for family composition, and assume income is shared within the family. We impose credit constraints during working life and retirement, permitting individuals to borrow only to fund education. This means that benefits may serve a useful role helping individuals transfer resources across periods of life.

Retirement is the same as working life, except that individuals no longer face a labour supply choice. Education decisions are made at age 17, before working life starts. Individuals choose between three levels: basic (corresponding to GCSEs), intermediate (A-levels or post-compulsory vocational education) and higher (university). This decision is made by comparing the expected lifetime utility of each option. We don’t say any more here about education since the

¹Dealing with rent properly would require an endogenous housing choice, creating considerable additional complexity that would be too hard to deal with in the current setup. Our approach is a reasonable approximation, but does not match perfectly the fraction of renters by family type (because certain family types are concentrated at particular ages) nor does it capture the fact that there are fewer renters at older ages.

results below hold education fixed.

By way of summary, we now set out the factors that will drive redistribution and insurance in the model: the dimensions heterogeneity across individuals (which determine tax liabilities and benefit entitlements) and the sources of uncertainty that individuals face. After the education decision but before the start of working life, individuals are heterogeneous in terms of the following variables: assets, k , unobserved preference for work, $\theta(l)$, access to free childcare, d^{CC} and education s . During life, additional heterogeneity evolves in terms of rent, H , experience, e , productivity, v , partner characteristics (presence of partner, d^m , education level, s^m , labour supply l^m and productivity v^m) and child characteristics (presence of children, d^k and age of youngest child, a^k). Together with female labour supply, these sources of individual heterogeneity determine tax liabilities and benefit entitlements and, as a result, redistribution and insurance.

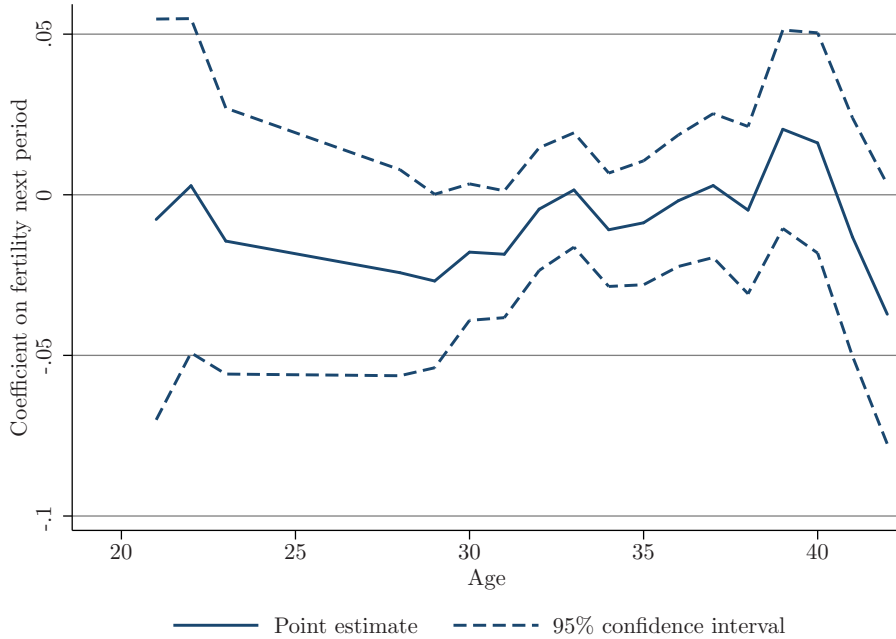
Insurance depends on the uncertainty individuals face, which is assumed to come from two sources: persistent productivity shocks that enter the wage process and stochastic changes to family composition (presence of partner and children) and partner employment and earnings. These are what generate the demand for insurance in the model. Insurance markets will be incomplete because some individual characteristics are not observable to third parties. Consequently, individuals will display precautionary behaviour by saving (to self-insure against future negative shocks) and will value any insurance provided by taxes and benefits (e.g. through wage subsidies or unemployment benefits).

If individuals partially foresee what we assume to be uncertainty then we will overstate the value of insurance provided by the tax and benefit system. Countering that, there are other sources of uncertainty that individuals face that we do not model. For example, Low et al (2010) emphasise the importance of employment risk, and we also do not include aggregate shocks. This is a big issue for structural and reduced form approaches alike, since they cannot possibly model all sources of uncertainty or all anticipated events that are relevant for decisions.

To assess the reasonableness of some of the assumptions in the in the model about the sources of uncertainty, we look for evidence of advance responses along the lines of Cunha et al. (2005). The idea is that if individuals know about something in the future we might expect to see them act upon that information. For example, women who know they are going to have a child in the near future may change their current education and labour market choices in anticipation. Given our model and the available data, we are fairly limited in the outcomes we can look for advance responses in; we use employment decisions. To begin with, we consider fertility and ask whether having a child next period can predict current employment decisions. Not conditioning on anything else, it is a strong predictor: among women without children this period, those who will have a child next period are 11 percentage points more likely to be working this period. But as soon as we start conditioning on the information the women does hold in the current period, the predictive power of fertility all but disappears. Figure 1 plots the effect of fertility next period on employment this period, controlling for

other information specified by the model as known today. The confidence bands straddle zero pretty much across all childbearing ages, providing no support for the advance information story.

Figure 1: Predictive power of fertility next period for current period employment



Notes: Graph presents results of a regression of current period employment on fertility next period, controlling nonparametrically for age and parametrically for education, experience, lagged employment and wage, presence of partner, partner employment status and partner earnings. The coefficient plotted corresponds to that on fertility next period.

In principle, the same exercise could be conducted for the other sources of uncertainty specified by the model: partner arrival and departure and wage shocks. Unfortunately, it seems hard to justify regressions using partner arrival or departure next period as an explanatory variable, since reverse causality seems likely to be a considerable issue. For example, women who are employed this period may make more attractive partners next period.

For completeness, we conclude this section by describing briefly some of the simplifications implied by our modelling framework. First, we exclude the value of public services. In other words, the results we present are for redistribution and insurance provided by the tax and benefit system, not for government activity as a whole. To see that this is an important omission, think about redistribution and insurance implicitly provided by the health service. It is hard to find a satisfactory way to estimate the value derived from each public service by each family, but attempting to derive such valuations would be a useful extension to this paper.

Since our aim is to understand the effects of the current tax system across the lifecycle, we model individuals as being subject to the same tax system throughout life. We ignore retirement, focusing instead on redistribution and insurance across working life. We do not model the response of wages and prices to taxes and benefits. This effectively means that we’re assuming that the incidence of taxes and benefits is fully on the individuals who pay the taxes or receive the benefits. Viewing the tax and benefit system from a lifecycle perspective raises the much bigger question of the intergenerational impact of the tax and benefit system. Almost all tax and benefit changes will have a differential effect across different age cohorts. Again, for simplicity, we do not address this.

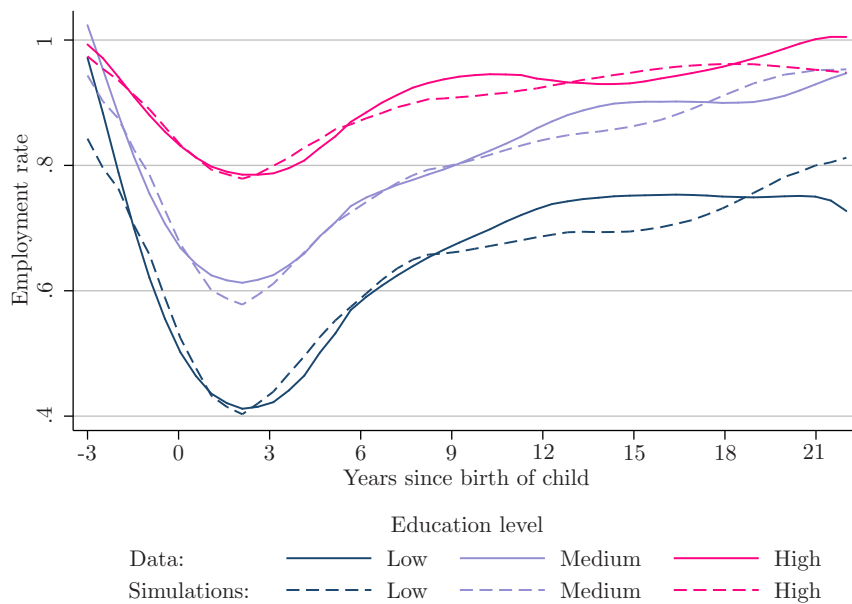
4.2 Data, estimation, validation and simulation

We estimate the model using data from the British Household Panel Survey (BHPS). Our dataset is an unbalanced panel of around 5,300 females and their families over 16 waves (1991-2006). Time effects are removed from wages. Estimation is done in three steps, summarised here and described in detail in Blundell et al. (2013). First, we use values from the literature to set the interest rate, discount rate and intertemporal preference parameter. Second, we estimate some parameters outside the model. These include family transitions, childcare costs, and male employment and wage processes). Finally, we estimate all the remaining parameters inside the model using the method of simulated moments (MSM). The parameters we estimate using MSM include those for the female wage equation, experience accumulation, preference parameters including taste for employment, unobserved heterogeneity and cost of education. We validated the model against non-structural estimates of the effects of recent UK tax reforms (see Blundell et al. (2013) for details). Table 1 sets out the values of some of the parameters likely to be important for our estimates of redistribution and insurance, including many of the parameter governing the female wage process. A complete list of parameter estimates is given in Blundell et al. (2013).

Figures 2 and 3 show that the model is able to capture the key features of female employment and wage dynamics. Figure 2 plots the female employment rate against years since birth of child, split by education level. The model matches the difference in employment rates across education levels and also the dip in employment around the birth of the child that is sharper for lower education levels. Figure 3 plots female log wages against age for different percentiles of the distribution, again split by education level. While there are some differences, in general the model does a good job of matching the age gradients that vary by education level and the dispersion across different percentiles.

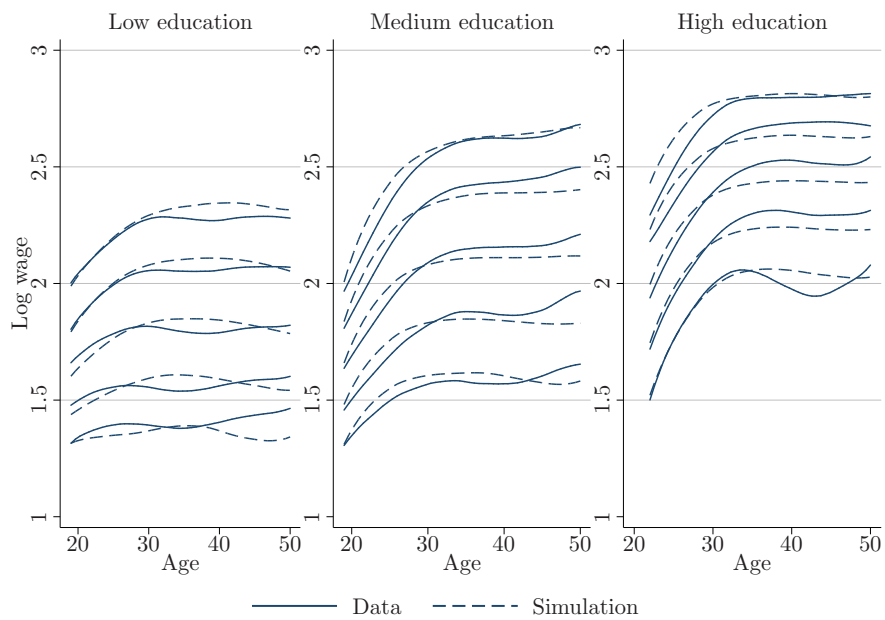
Simulated lifecycles form the basis for the analysis presented below. Initial conditions (at age 17) are taken from the BHPS data for a representative sample of 562 individuals. For each individual, we draw stochastic lifecycle profiles for the exogenous processes (productivity and family composition) and solve the decision problem at each age. The result is a set of individual lifecycle

Figure 2: Model fit: employment rate by education level



Notes: Graph compares BHPS data with data simulated by the model. Series smoothed.

Figure 3: Model fit: log wage rate by education level for various percentiles (10, 25, 50, 75 and 90)



Notes: Graph compares BHPS data with data simulated by the model. Series smoothed.

Table 1: Parameter estimates for key model parameters

Parameter	Value
Interest rate (r)	0.015
Discount rate (β)	0.980
CRRA preference parameter (γ)	1.560
Market wage rate (W_s)	
Basic education	4.496
Intermediate education	4.895
Higher education	6.261
Return to experience (α)	
Basic education	0.133
Intermediate education	0.240
Higher education	0.271
Standard error of innovation to productivity (σ_{ϵ_s})	
Basic education	0.123
Intermediate education	0.126
Higher education	0.105
Persistence of productivity shock (ρ_s)	
Basic education	0.954
Intermediate education	0.952
Higher education	0.922

profiles for each of the exogenous and endogenous variables in the model (e.g. labour supply, consumption, assets, experience, education). We then use these simulations to provide initial conditions for simulations from the start of working life, replicating each 100 times.

We simulate under each of the baseline, reform and intermediate tax and benefit systems described above and use the results to calculate the value of each redistribution and insurance component. These intermediate tax and benefit systems require the model solution to be recomputed multiple times, reflecting the fact that the adjustments used to compute the intermediate systems vary across individuals, age and sometimes replication. Given the focus of the model on females, the population we simulate corresponds to families containing an adult female (i.e. single adult male families are excluded).

5 Tax and benefit reforms

We now provide a brief overview of the UK tax and benefit system and the reforms we will decompose in Section 6. Our focus is on personal taxes and benefits, so we exclude corporate, capital and indirect taxes. We also concentrate on employment earnings during working life, ignoring other forms of income such as self-employment profits, pensions and savings income. We don't consider disability and ignore the contributory nature of parts of the tax and benefit system

(i.e. payment of particular taxes accumulates future benefit entitlement) since this link is very weak in practice. A more comprehensive discussion of the UK system can be found in Adam and Browne (2011) and Jin and Phillips (2010).

5.1 UK tax and benefit system overview

The UK personal tax and benefit system comprises a small number of simple taxes (mostly levied on individuals), and a complex web of benefits and tax credits (usually means-tested at the family level).

Taken together, personal taxes on employee earnings operate through a system of tax-free allowances and income bands that are subject to a progressive schedule of tax rates. Most employed earners face a combined marginal tax rate on earnings of 32 per cent. The other main personal tax is an annual charge based on house value (i.e. not earnings-related).

Benefits and tax credits are in the process of being rationalised (see below). Nevertheless, support retains broadly the same structure under all of the systems we are considering. Workless families receive an income top-up to some subsistence level and potentially additional payments towards rent, council tax and the cost of children. The value of these benefits varies depending on family composition. Working families gradually lose entitlement to this support as their earnings rise, though until the current reforms there was some in-work support contingent on the number of hours worked (and also tapered away from richer families). We now describe the reforms of interest in some detail.

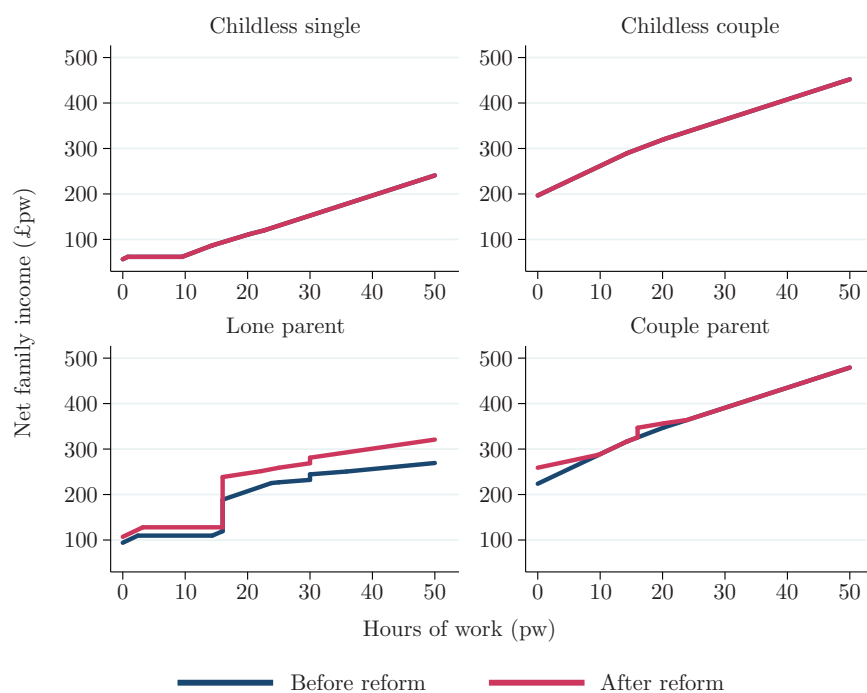
5.2 Working families tax credit reform (1999)

In October 1999, one in-work benefit (family credit, FC) was replaced by another one (working families tax credit, WFTC). Although the two benefits were structurally very similar, WFTC was considerably more generous than FC, for three main reasons. First, maximum awards were higher. For example, the maximum award for a lone parent working 16-29 hours with one child under 11 rose in real terms by 25% if there was no childcare expenditure and 162% with the maximum £135 a week childcare expenditure. Second, the threshold above which awards were tapered away was higher, rising in real terms by 10%.² And third, awards were tapered away more slowly (55% rather than 70%). The combined effect of these changes was to increase substantially awards for existing claimants and extend entitlement to new (richer) families.

At the same time, the generosity of out-of-work benefits (income support and income-based jobseeker's allowance) was increased for families with children (particularly younger children). We consider this as part of the same reform since in- and out-of-work benefit changes were designed alongside each other.

²Note, however, that the FC threshold was effectively much higher if the FC childcare disregard is taken into account. So, although the maximum award under FC for a family with childcare expenditure was much lower, it was tapered away from a considerably higher level. For families on the FC taper, this meant that the increase in generosity under WFTC wasn't nearly as big as implied by the increase in the maximum award.

Figure 4: Budget constraints before and after WFTC reform



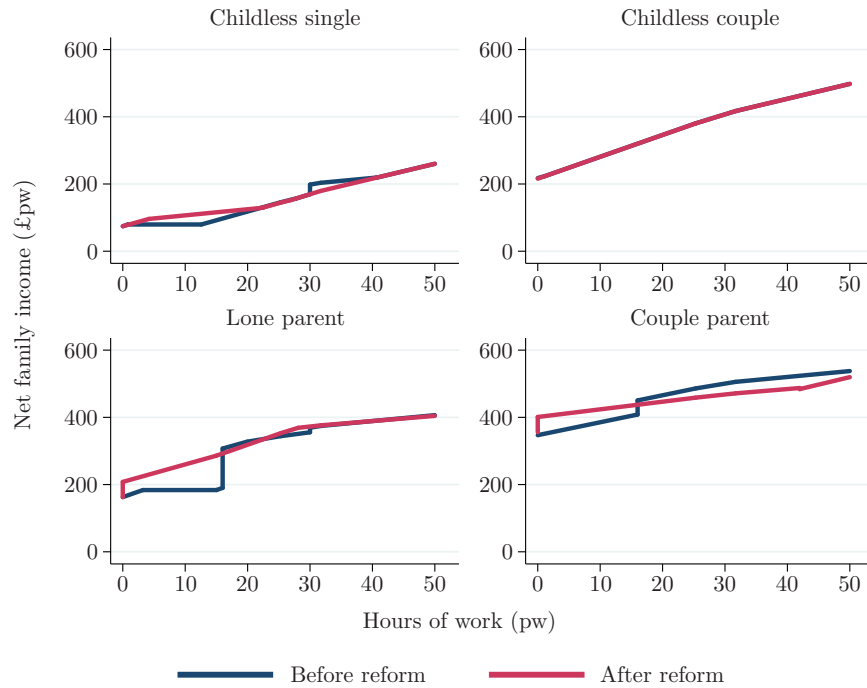
Notes: All individuals are assumed to earn £6 per hour. Families with children are assumed to have one child aged four and spend £50 per week on childcare. Rent is zero in all cases. Tax and benefit systems updated to 2013 levels.

Taken together, the WFTC and out-of-work benefit reforms were substantial giveaways: based on the simulations reported below, it amounted to a giveaway of £123 per year on average for each female. The main family types to gain were lone parents. Couple families were incentivised to have one earner rather than two.

Figure 4 gives example budget constraints for a variety of different family types under the April 2002 tax and benefit system before and after the WFTC (and IS) reform. It is clear from this figure that low paid lone parents were the main beneficiaries of the reform. Those without children did not gain.

5.3 Universal credit reform (2013)

Figure 5: Budget constraints before and after UC reform



Notes: All individuals are assumed to earn £6 per hour. Families with children are assumed to have one child aged four and spend £50 per week on childcare. Rent is zero in all cases. Tax and benefit systems updated to 2013 levels.

The introduction of universal credit (UC) over a five-year period starting from April 2013 amounts to perhaps the most radical restructuring of the benefits system since the 1940s. UC will replace six of the seven main means-tested benefits and tax credits for working-age individuals, and implies substantial

changes in the level of support for, and work incentives among, some types of family.

Entitlements for those with no other income or assets are the same as under the previous system. But there are two key changes that affect entitlements in work. First, UC awards are withdrawn at a single unified rate (65%) as earnings increase rather than being subject to multiple overlapping tapers as under the previous system. This reduces the maximum marginal effective tax rate an individual can face quite considerably (e.g. it falls below that created by tax credits and HB together). Second, there are no longer any jumps in the budget constraint when an individual works a certain number of hours each week (e.g. 16 and 30) as there were under the previous system.

Unlike the WFTC reform, UC was designed to be broadly revenue neutral overall, meaning that (ignoring transitional protection) some family types gain at the expense of others. The main family type to gain on average is one-earner couples with children. Family types that lose on average include workless households and lone parent households. Figure 5 gives example budget constraints for a variety of different family types under the April 2015 pre-UC and post-UC tax and benefit systems. These demonstrate some of the key features of the reform, including the changes in taper rates and elimination of jumps at given hours points. And, as they suggest, childless couple families were least likely to be affected by the reforms.

In addition to these changes to awards financial work incentives, UC involves other important non-financial changes. These include the administrative simplification associated with combining six benefits into one, and the greater conditionality (additional requirements placed on claimants to find work and increase their earnings) being introduced. We don't model either of these.

6 Results

In this section, we set out the decomposition results, first for the WFTC reform, then for the UC reform and finally for a hypothetical reform from a revenue-neutral flat-rate baseline system to the April 2015 system including the UC reform.

6.1 WFTC

We begin by holding behaviour fixed at its baseline level, by which we mean hours worked and the proportion of total resources consumed are both unchanged.³ The baseline system we use is the April 2002 system but with family credit (WFTC's predecessor) and income support uprated in line with default uprating rules rather than the reforms that actually happened.

Overall, the WFTC (and income support) reform amounted to an average net giveaway of £123 per year for each female, equivalent to 1.3 per cent of

³We can't hold the level of consumption fixed because individuals' budget constraints may no longer be satisfied.

net annual income (2). But this giveaway was not evenly distributed across the population. As the top left graph in Figure 6 shows, from a cross-sectional viewpoint, it strongly favoured lower deciles, particularly the bottom two deciles.

A common measure of the progressivity of a tax or benefit reform is the Kakwani index of tax progressivity. This index describes the disproportionality of the reform relative to pre-reform incomes. A positive index indicates that the reform is progressive. For a benefit reform that has no losers, the maximum value for the Kakwani index in the absence of behavioural responses is 2 (there is no such constraint once behavioural responses are allowed for). Treating the simulated data as a cross-section, the Kakwani index for the WFTC reform (allowing for behavioural responses) is 0.75, a highly progressive reform.

But things look rather different from the perspective of the start of working life. Table 2 decomposes the total size of the WFTC reform into the sum of a giveaway term, and four redistribution and insurance components, all in expressed annual per-female terms. The redistribution and insurance components are all zero in the first two lines by definition. For example, redistribution between individuals involves taking from one individual and giving to another, exactly offsetting at the aggregate level. (We include them for expositional purposes only). In order to assess the size of the reform, it is therefore necessary to work in terms of something that will not exactly offset; we use absolute values. On this basis, the penultimate line of the table shows that the largest components of the reform were the giveaway, insurance between and insurance within. These come in at -1.3, 0.8 and 1.0 per cent of net income respectively (the figure for the giveaway is negative because it is a reduction in total net tax).⁴

Table 2: Decomposition of WFTC reform

	Mean giveaway	Redistribution		Insurance	
		Between	Within	Between	Within
Change (£)	123	0	0	0	0
Share of net income	0.013	0.000	0.000	0.000	0.000
Change in abs value (£)	-123	44	43	72	87
Share of net income	-0.013	0.004	0.005	0.008	0.010
Share of total	-0.999	0.361	0.349	0.583	0.706

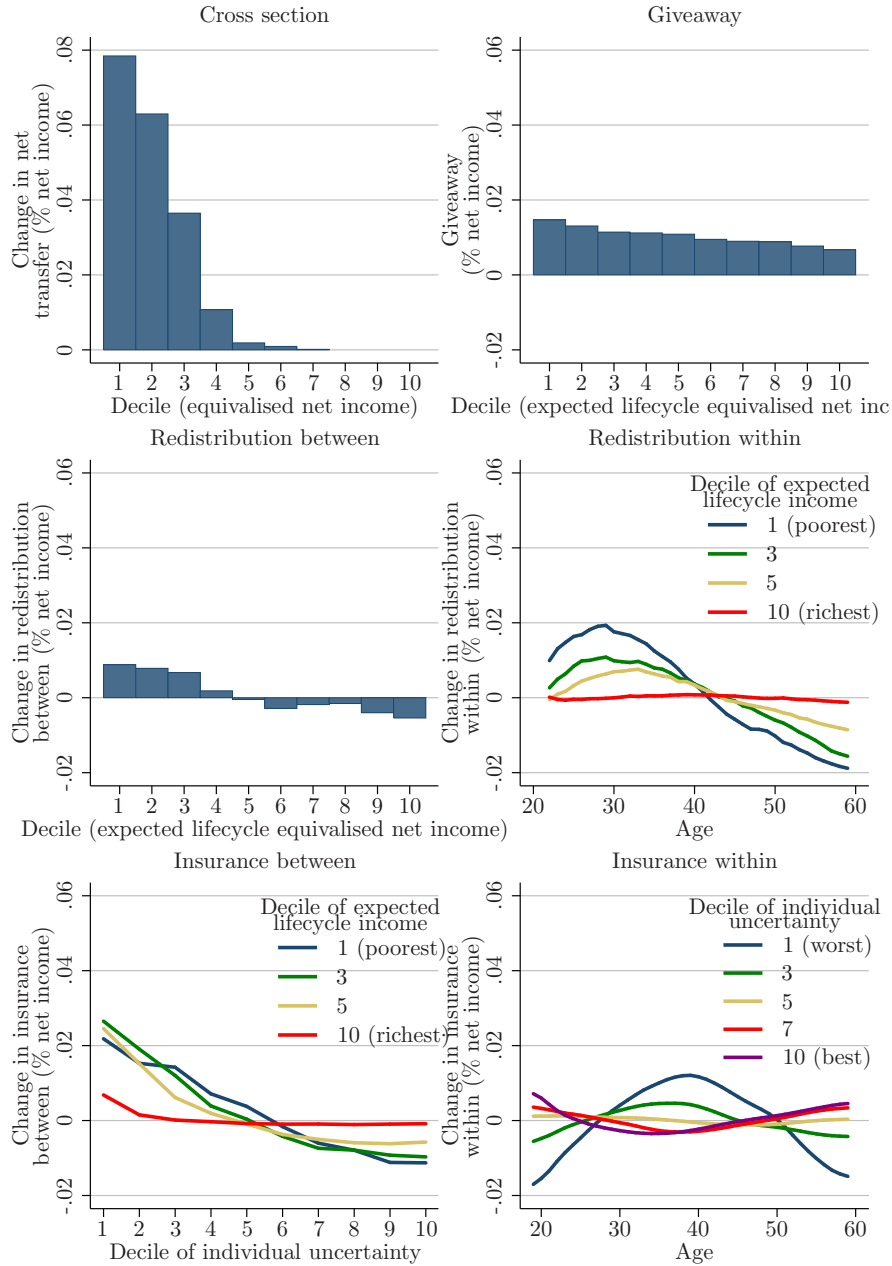
Monetary figures expressed in annual per-female terms.

We now consider how each of these components varies across individuals. Figure 6 shows the impact of all five components. The top right graph gives the distributional impact of the lump-sum giveaway component, where individuals are put into deciles on the basis of expected lifecycle income from the perspective of the start of working life. Not surprisingly, expressed as a percentage of annual net income, the giveaway declines gradually in importance for higher deciles.

The centre left graph in Figure 6 shows the distributional impact of the

⁴Note that redistribution and insurance components partially offset each other, so the sum of the components exceeds the total change in families' budgets.

Figure 6: Cross-sectional impact and decomposition of WFTC reform



between-individual redistributive component, which ranges from 0.9 per cent of net income for females in the bottom lifecycle income decile to -0.6 per cent at the top. This is noticeably less progressive than the cross-sectional graph in the top left corner. Indeed, if we calculate the Kakwani progressivity index for the combined effect of the lump-sum giveaway and redistribution between components, we get a value of 0.39, much lower than the 0.75 we got when viewing the reform in the cross section.

As indicated by Table 2, the change in redistribution within individuals (across periods of life) is as large as the change in redistribution between individuals. We can see from the centre right graph in Figure 6 what form this takes. There are two main things to draw from this graph. First, in general, the tax and benefit system tends to move resources forward in time, i.e. from later in life to earlier in life. Second, this is particularly pronounced for poorer deciles. For the richest decile, there is very little redistribution across periods of life at all.

The bottom left graph in Figure 6 shows the pattern of how between-individual insurance changes in response to the reform. In previous graphs, we have been arranging women along the horizontal axis by expected lifecycle income. In this figure, we order by the uncertainty they face: from the worst 10 per cent of outcomes for each woman on the left to the best 10 per cent of outcomes for each woman on the right.⁵ In general, between-individual insurance transfers are positive for the bottom 50 per cent of realisations of uncertainty. What is striking, however, is the pattern by expected lifecycle income decile.

While all deciles gain most when uncertain outcomes are worst, deciles 3 and 5 actually do better from the reform than the poorest decile right at the bottom. It is only for mid-range outcomes of uncertainty that the bottom expected lifecycle income decile does better. This is probably because the reform helped low-paid working families most, very few individuals in the bottom expected lifecycle income decile and with the worst uncertain outcomes fall into that group.

Finally, the bottom right graph in Figure 6 shows the pattern of within-individual insurance across the lifecycle. In this case, the horizontal axis is age and each line is a selected decile of individual uncertainty. In general, we see that worse realisations of shocks result in more insurance towards the middle of working life.

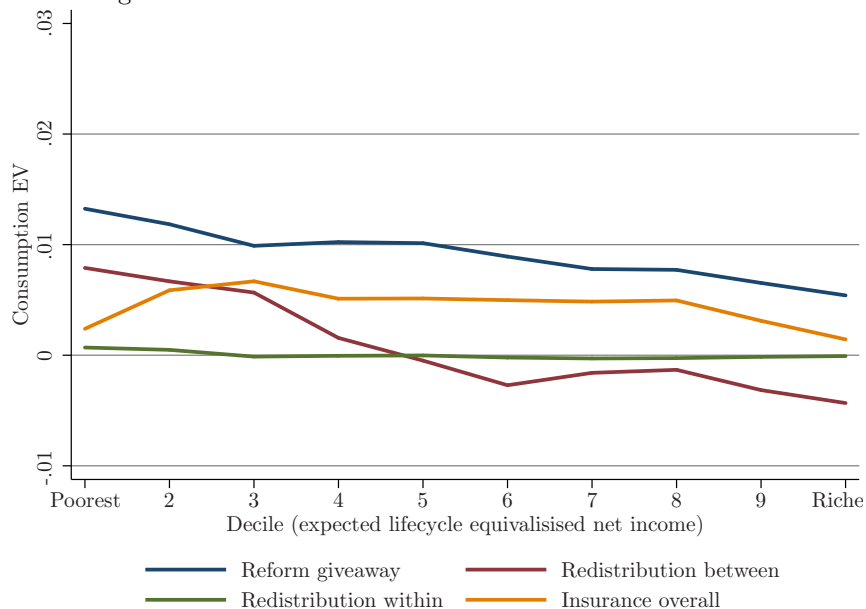
All results presented so far have related to the quantity of resources moved around by the tax and benefit system; we now consider utility valuations. As described above, we calculate utility valuations as average consumption equivalent variations, i.e. the proportional increase in consumption that individuals require to be as well off under the baseline system as they are with the reform. We have been unable to come up with an adequate way of separating the two insurance components, so here report the combined value of the insurance components.

Figure 7 shows the utility valuations for the WFTC reform, split by decile of expected lifecycle equivalised net income. As with the preceding figures,

⁵Best and worst outcomes are defined on the basis of equivalised lifecycle net income

behaviour is held fixed (we consider behavioural responses below). As expected, the value of the reform giveaway is positive and declines with the decile, as does the value of between-individual redistribution. For the latter, the value is positive for the bottom half of the distribution and negative for the top half. Within-individual redistribution is worthless on average for all but the bottom decile, and even there its value corresponds to just 0.1 per cent of consumption. This is despite the fact that the model imposes fairly stringent borrowing constraints (borrowing is only allowed to fund education).⁶ The value of insurance is positive for all deciles but particularly in the middle of the income distribution. Taking all components together apart from the giveaway, the bottom eight deciles all gain on average, though for deciles six to eight the gain is modest.

Figure 7: Consumption-equivalent utility valuations of WFTC reform components holding behaviour fixed



We now consider what effect allowing behavioural responses in labour supply and consumption has on the foregoing results. Table 3 shows that we estimate an overall employment impact of -0.8 per cent, driven by offsetting effects for

⁶The low value attached to within-individual redistribution unlikely to be a consequence of behaviour being held fixed. We take fixed behaviour to mean that labour supply and the share of consumption out of total resources (not its level) are both unchanged. Therefore, to the extent that the reform relaxes borrowing constraints, the consumption of constrained individuals (who consume all their resources) will increase by the full amount of the relaxation. The reason we hold the consumption share rather than the level fixed is that holding the level fixed might violate individuals' budget constraints.

lone mothers (6.4 per cent) and mothers in couples (-4.6 per cent). Almost all of the increase in employment for lone mothers is in part-time employment; for mothers in couples the employment fall is split roughly equally between part- and full-time employment. These estimates are not directly comparable with previous results in the literature (e.g. because we impose the reform tax and transfer system across the whole adult life rather than on a cross-section at a point in time). Nevertheless, they exhibit broadly the same patterns as in earlier work (e.g. Blundell et al. (2005)).

Table 3: Estimated impact of WFTC on employment

	Employment	Part-time employment
Overall	-0.008	0.000
Childless single	0.000	0.000
Lone mother	0.064	0.060
Childless couple	-0.001	0.003
Couple with children	-0.046	-0.025

Table 4 shows the decomposition of the WFTC reform allowing for behavioural responses. Relative to Table 2, changes in behaviour increase the size of the average giveaway by just over 30 per cent. This is due partly to the overall decline in employment and partly because the fiscal saving from lone mothers moving into part-time work is likely to be modest compared the the additional fiscal burden from mothers in couples moving out of work. The table also shows that changes in behaviour increase the size of the within-individual redistribution component. Nevertheless, the corresponding graphs are qualitatively similar so we don't present them here.

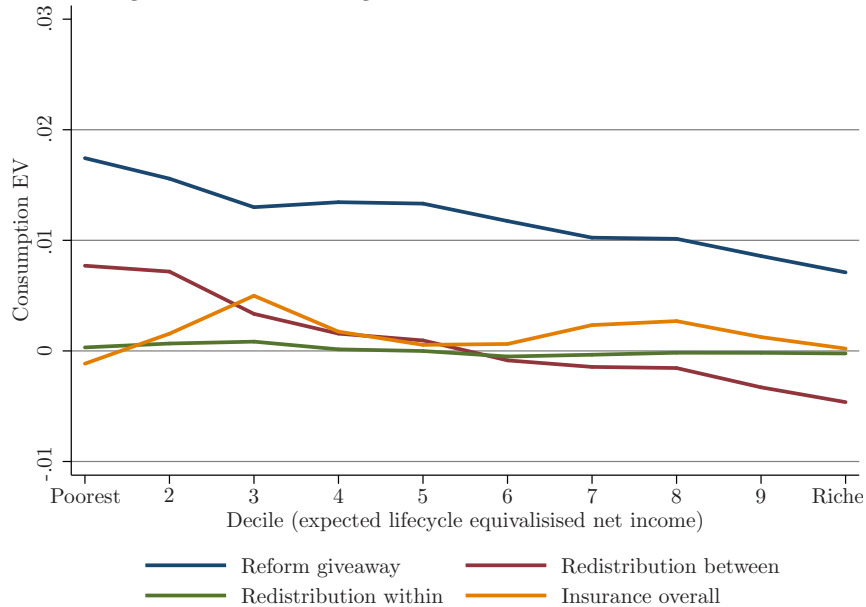
Table 4: Decomposition of WFTC reform allowing for behavioural response

	Mean giveaway	Redistribution		Insurance	
		Between	Within	Between	Within
Change (£)	162	0	0	0	0
Share of net income	0.017	0.000	0.000	0.000	0.000
Change in abs value (£)	-162	45	75	64	72
Share of net income	-0.017	0.005	0.008	0.006	0.004
Share of total	-1.754	0.485	0.807	0.687	0.774

Monetary figures expressed in annual per-female terms.

We now consider how the utility valuations change once we allow for behavioural responses. This is shown in Figure 8. While the broad pattern is the same, behavioural responses tend to increase the value of the giveaway and reduce the value of insurance.

Figure 8: Consumption-equivalent utility valuations of WFTC reform components allowing behaviour to change



6.2 UC

We now turn to analysis of the UC reform. The baseline against which we judge the effects of the reform is the projected April 2015 system but without UC and with the benefits that UC replaces uprated in line with default uprating rules.

In contrast to WFTC, UC was broadly revenue neutral, amounting to an annual net giveaway of £33 per female (around 0.3 per cent of net income) holding behaviour fixed. Despite being close to revenue neutral, there are still some important distributional consequences viewed from a cross-sectional perspective. This is shown in the top left panel in 9. In particular, there are strong losses (lower benefits) in the bottom decile amounting to 3 per cent of net income. This is driven by the reduction in generosity of the benefits system to lone parents under 25.

Table 5 holds behaviour fixed and decomposes the total size of the UC reform in the same way as Table 2 did for WFTC. The reform is a small takeaway (0.3 per cent of net income) and there are modest declines in the size of all the other components.

The remaining graphs in Figure 9 show how the components of the decomposition vary across individuals. The top right graph shows the giveaway component, which is negligible. The centre left graph shows that there was little change in the between-individual redistribution component; what there was, tended to favour those in the richer deciles at the expense of those in poorer

Figure 9: Cross-sectional impact and decomposition of UC reform

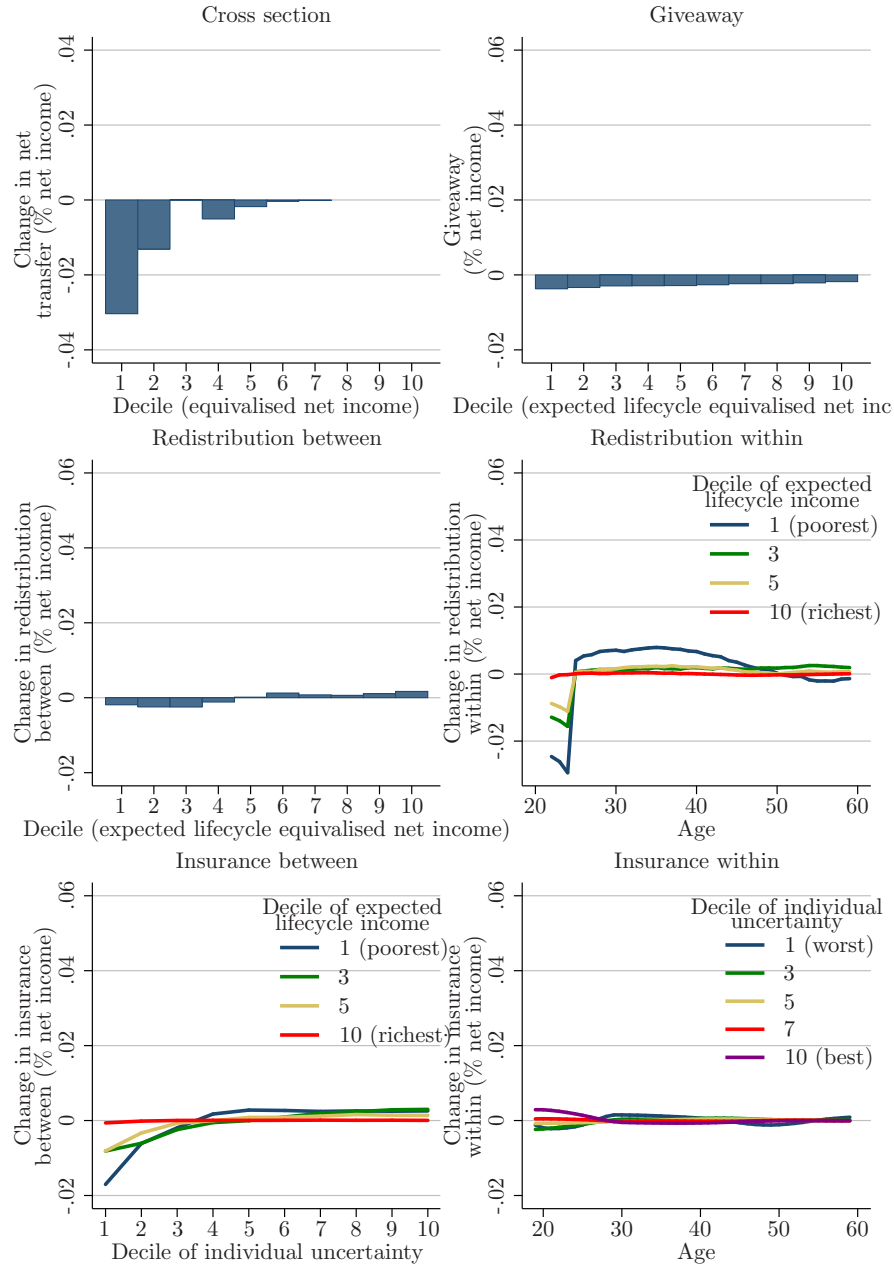


Table 5: Decomposition of UC reform

	Mean giveaway	Redistribution		Insurance	
		Between	Within	Between	Within
Change (£)	-33	0	0	0	0
Share of net income	-0.003	0.000	0.000	0.000	0.000
Change in abs value (£)	33	-15	-7	-15	-25
Share of net income	0.003	-0.001	-0.001	-0.002	-0.003
Share of total	-1.096	0.506	0.249	0.514	0.827

Monetary figures expressed in annual per-female terms.

deciles. There is quite a clear pattern for within-individual redistribution (centre right graph), with resources transferred discretely from before age 25 to after age 25. As discussed earlier, this is the consequence of changes in entitlements for lone parents under age 25.

There are not many substantial changes in terms of between-individual insurance (bottom left graph), though there is a notable reduction in the amount of insurance against the worst 10 per cent of shocks for the bottom few deciles. In terms of within-individual uncertainty (bottom right graph), most of the main changes happen before age 30, with lower deciles receiving relatively less insurance than later in life and higher deciles relatively more insurance.

Figure 10 shows average utility valuations of the UC reform, holding behaviour fixed and split by decile of expected lifecycle equivalised income. While none of the components by themselves is large, the combined effect across all components (excluding the giveaway component) is negative for the bottom six deciles, and is less than -0.5 per cent for both of the bottom two deciles.

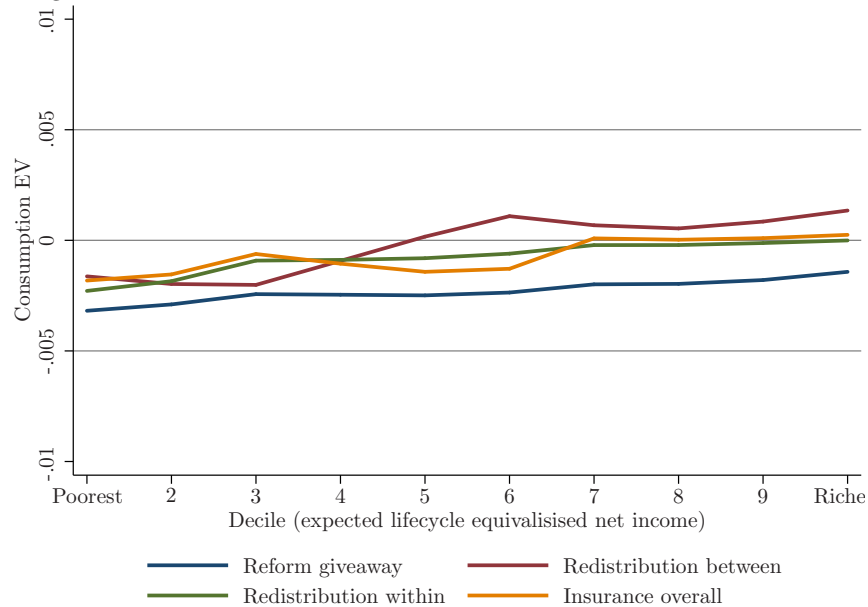
We now consider how results change once we allow for behavioural responses. Table 6 sets out estimated employment effects of the UC reform. Overall there is a 0.8 per cent increase in the employment rate and a 1.3 per cent increase in part-time employment. Lone parents are affected much more than the other family types, experiencing a 4.8 per cent increase in employment and a 6.7 per cent increase in part-time employment. In large part, these are a consequence of entitlements being less generous for lone parents under age 25.

Table 6: Estimated impact of UC on employment

	Employment	Part-time employment
Overall	0.008	0.013
Childless single	0.003	0.003
Lone mother	0.048	0.067
Childless couple	0.006	0.007
Couple with children	-0.001	0.006

Table 7 sets out what happens to the reform components once we allow

Figure 10: Consumption-equivalent utility valuations of UC reform components holding behaviour fixed



behavioural responses. The size of the takeaway falls; other components exhibit modest changes.

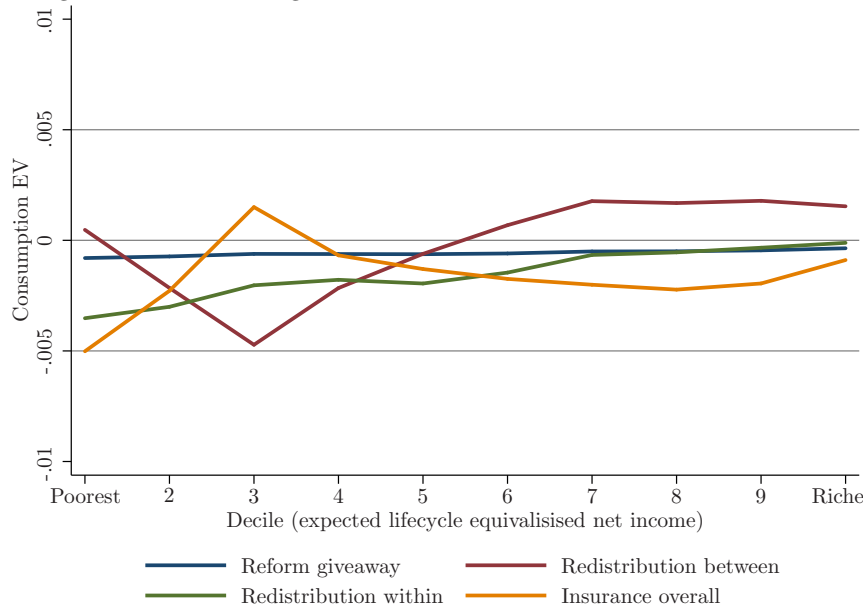
Table 7: Decomposition of UC reform allowing for behavioural response

	Mean giveaway	Redistribution		Insurance	
		Between	Within	Between	Within
Change (£)	-8	0	0	0	0
Share of net income	-0.001	0.000	0.000	0.000	0.000
Change in abs value (£)	8	-27	-17	-6	-21
Share of net income	0.001	-0.002	-0.002	-0.001	-0.006
Share of total	-0.133	0.430	0.278	0.088	0.337

Monetary figures expressed in annual per-female terms.

Figure 11 shows average utility valuations of the UC reform. A comparison with Figure 10 shows that allowing for behavioural responses make the value of the giveaway component somewhat less negative, but do the opposite for the redistribution within component. It also increases the value of insurance and decreases the value of between-individual redistribution, particularly in the third decile.

Figure 11: Consumption-equivalent utility valuations of UC reform components allowing behaviour to change



6.3 UC relative to a flat-rate baseline

Finally, we set out the decomposition for the UC system relative to a hypothetical flat-rate baseline that raises the same net revenue from females. The aim of considering such a reform is to determine what a modern tax system achieves through its complex web of taxes and benefits relative to a very simple alternative. The flat-rate baseline we use gives each individual a lump-sum transfer equal to the main rate of UC for single adults and subjects individuals to a marginal tax rate chosen to ensure revenue-neutrality under fixed behaviour. The marginal tax rate that achieves this is 38.5 per cent.

The top left graph in Figure 12 shows that, from a cross-sectional perspective and holding behaviour fixed, the reform primarily benefits the bottom decile at the expense of deciles three to seven.

Table 8 sets out the monetary decomposition for fixed behaviour. Not surprisingly, the size of entries in the change in absolute value row is much bigger than it was under either of the previous reforms. This reform does considerably less between-individual redistribution (corresponding to 0.9 per cent of net income) and considerably more within-individual redistribution and insurance (1.3 and 4.2 per cent of net income respectively).

Figure 12 sets out each of the components of the decomposition. Note these graphs have an enlarged y-axis scale, so many of the changes observed here are considerably larger than those for the previous two reforms.

Figure 12: Cross-sectional impact and decomposition of UC reform from flat-rate baseline

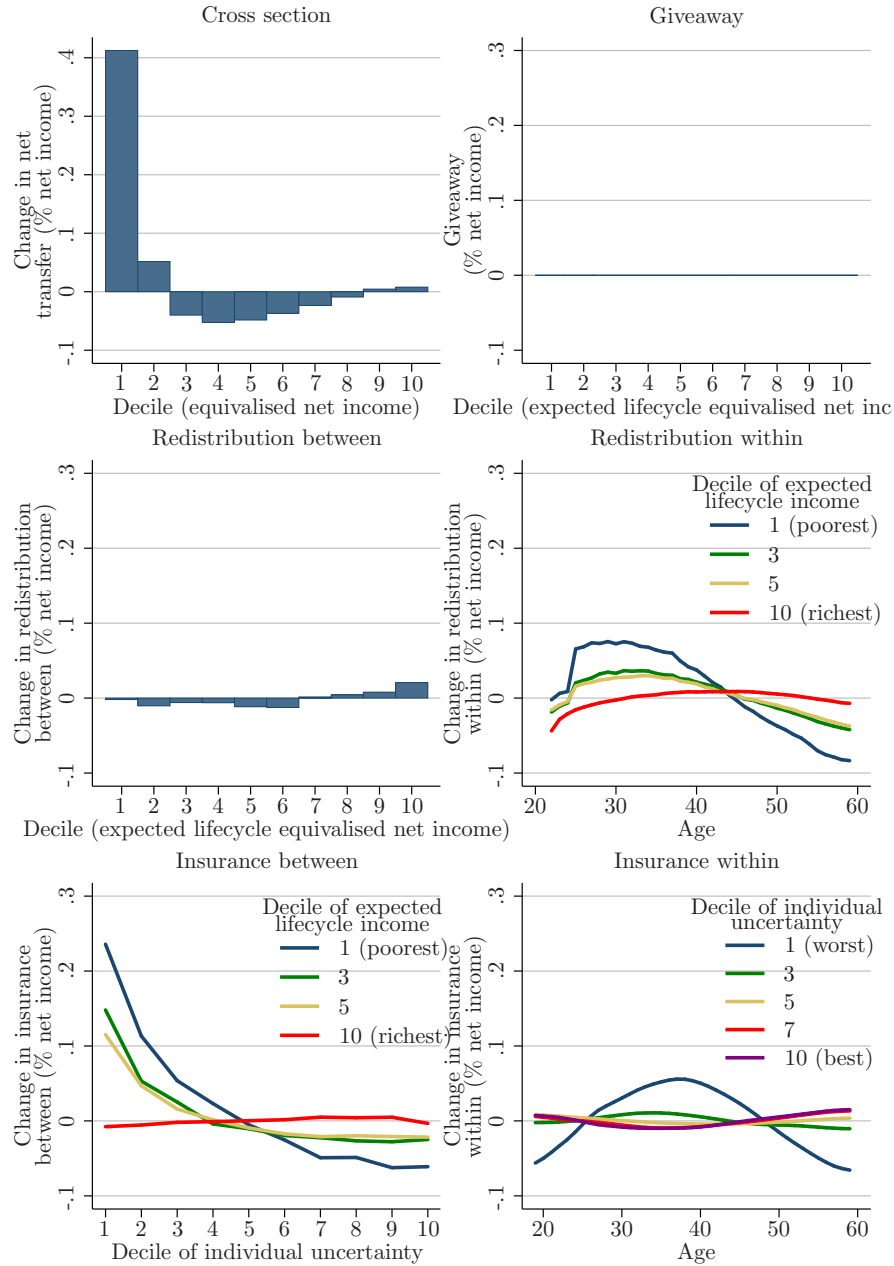


Table 8: Decomposition of UC reform relative to flat-rate baseline

	Mean giveaway	Redistribution		Insurance	
		Between	Within	Between	Within
Change (£)	0	0	0	0	0
Share of net income	0.000	0.000	0.000	0.000	0.000
Change in abs value (£)	0	-104	94	-25	220
Share of net income	0.000	-0.009	0.013	0.006	0.042
Share of total	0.000	-0.560	0.507	-0.134	1.187

Monetary figures expressed in annual per-female terms.

The top right graph shows the reform giveaway is zero because it is revenue neutral. Perhaps somewhat surprisingly, the centre left graph shows that reform achieves approximately the same amount of redistribution between individuals as the flat-rate baseline. If anything, it implies a small amount of redistribution from the poorest deciles to the richest ones. The centre right graph shows that the UC system does considerably more within-individual redistribution towards the main childbearing years. This is particularly true for lower deciles. But the key difference of the UC reform is summarised by the bottom two graphs: the UC system provides a much more substantial insurance safety net against the worst outcomes than the flat-rate system. This is true for both between-individual insurance (where the bottom couple of deciles of individual uncertainty are insured against considerably more than under the baseline) and within-individual insurance (where resources are targeted towards the main child-rearing ages).

Utility valuations holding behaviour fixed are in Figure 13. The real stand-out feature of this graph is the very substantial value attached to the insurance component, particularly towards the bottom of the income distribution. Indeed, for the poorest two deciles, the additional insurance provided by the UC system relative to the flat-rate baseline is equivalent to an increase in consumption exceeding 5 per cent.

We now turn to behavioural responses. Table 9 sets out estimated employment effects of the UC reform relative to a flat-rate baseline. Overall there is a 3.8 per cent increase in employment and a 4.4 per cent increase in part-time employment. Lone parents move out of employment (a 4.8 per cent fall) and switch strongly in favour of part-time employment (a 20.6 per cent rise). The employment rate for women in childless couples increases by 8.9 per cent, most of which is full-time employment, and the employment rate for mothers in couples rises by 7.6 per cent, most of which is part-time.

Table 10 sets out what happens to the reform components once we allow behavioural responses. From the first column, what was a revenue-neutral reform under no behavioural response has become a substantial net giveaway, corresponding to 2.9 per cent of net income. There is a small fall (to 0.3 per cent of net income) in the volume of between-individual redistribution transactions.

Figure 13: Consumption-equivalent utility valuations of UC reform components relative to fixed-rate baseline holding behaviour fixed

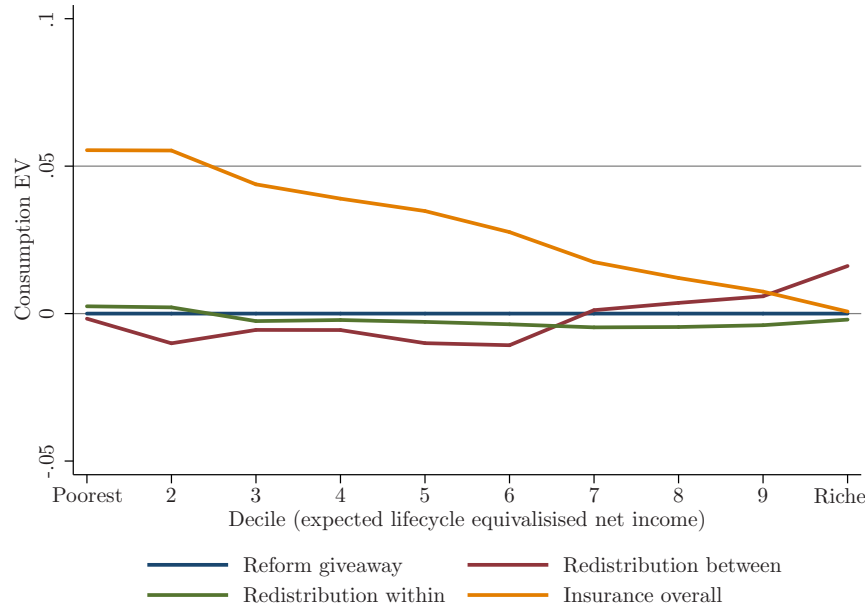


Table 9: Estimated impact of UC reform relative to a flat-rate baseline on employment

	Employment	Part-time employment
Overall	0.038	0.044
Childless single	-0.019	-0.000
Lone mother	-0.048	0.206
Childless couple	0.089	-0.015
Couple with children	0.076	0.065

The remaining three components experience substantial increases, especially the within-person insurance term, which rises by 5.5 per cent of net income.

Table 10: Decomposition of UC reform relative to flat-rate baseline allowing for behavioural response

	Mean giveaway	Redistribution		Insurance	
		Between	Within	Between	Within
Change (£)	282	0	0	0	0
Share of net income	0.029	0.000	0.000	0.000	0.000
Change in abs value (£)	-282	-38	172	148	449
Share of net income	-0.029	-0.003	0.018	0.022	0.055
Share of total	-0.629	-0.085	0.382	0.331	1.001

Monetary figures expressed in annual per-female terms.

Figure 14 shows average utility valuations of the reform, allowing behaviour to change. Relative to fixed behaviour, the size of the insurance component is reduced across the income distribution but the size of the reform giveaway component increased.

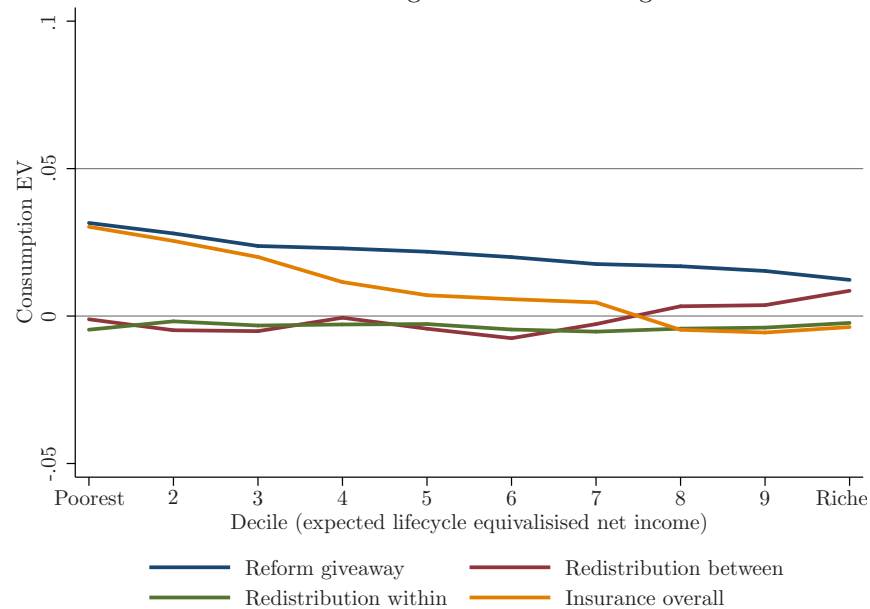
7 Conclusion

In this paper we set out to understand the roles that taxes and transfers play in redistributing resources and providing insurance across individuals and across the lifecycle. We decomposed net transfers into five components: a giveaway term and terms corresponding to between- and within-individual redistribution and between- and within-individual insurance. On the basis of this decomposition we have been able to demonstrate what the tax and transfer system achieves from a lifecycle perspective and why it is valuable.

We compared the 2015 UK tax and transfer system relative to a flat-rate baseline, showing what is achieved by the complex tax and welfare entitlement rules in a modern economy. We also assessed the impact of two important UK benefit reforms—the working families’ tax credit (WFTC) reform of 1999 and the universal credit (UC) reform that began in 2013. We found that insurance against wage and family composition shocks is substantial and highly valued by individuals. Within-individual redistribution (i.e. across periods of life) is generally of little value even in the presence of strict borrowing constraints. Behavioural responses tend to increase the size of reform giveaways at the expense of the other components.

This is an important first step towards a much bigger agenda trying to understand what effects taxes and benefits have when viewed from a lifecycle perspective rather than the standard static outlook—an important shift in perspective if we are to have a more complete understanding of the tax and benefit system.

Figure 14: Consumption-equivalent utility valuations of UC reform components relative to fixed-rate baseline allowing behaviour to change



References

- Stuart Adam and James Browne. A survey of the uk tax system. 2011.
- J Adda, M Costa Dias, C Meghir, and B Sianesi. Labour market programmes and labour market outcomes: a study of the swedish active labour market interventions. 2007.
- Charlotte Bartels. Redistribution and insurance in the german welfare state. 2011.
- Anders Björklund and Mårten Palme. *Income redistribution within the life cycle versus between individuals: Empirical evidence using Swedish panel data*. Stockholm School of Economics, the Economic Research Inst., 1997.
- Richard Blundell, Mike Brewer, and Andrew Shephard. Evaluating the labour market impact of working families' tax credit using difference-in-differences. 2005.
- Richard Blundell, Monica Costa Dias, Costas Meghir, and Jonathan M Shaw. Female labour supply, human capital and welfare reform. Technical report, National Bureau of Economic Research, 2013.

- A Lans Bovenberg, Martin Ino Hansen, and Peter Birch Sørensen. Individual savings accounts for social insurance: rationale and alternative designs. *International Tax and Public Finance*, 15(1):67–86, 2008.
- Mike Brewer, Alan Duncan, Andrew Shephard, and María José Suarez. Did working families’ tax credit work? the impact of in-work support on labour supply in great britain. *Labour Economics*, 13(6):699–720, 2006.
- Raj Chetty. Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, 116(2):173–234, 2008.
- Flavio Cunha, James Heckman, and Salvador Navarro. Separating uncertainty from heterogeneity in life cycle earnings. *Oxford Economic Papers*, 57(2): 191–261, 2005.
- Zvi Eckstein and Kenneth I Wolpin. Dynamic labour force participation of married women and endogenous work experience. *The Review of Economic Studies*, 56(3):375–390, 1989.
- Zvi Eckstein and Kenneth I Wolpin. Why youths drop out of high school: The impact of preferences, opportunities, and abilities. *Econometrica*, 67(6): 1295–1339, 1999.
- Nada Eissa and Hilary W Hoynes. Behavioral responses to taxes: Lessons from the eitc and labor supply. In *Tax Policy and the Economy, Volume 20*, pages 73–110. The MIT Press, 2006.
- Nada Eissa and Jeffrey B Liebman. Labor supply response to the earned income tax credit. *The Quarterly Journal of Economics*, 111(2):605–637, 1996.
- Jane Falkingham and Ann Harding. Poverty alleviation vs social insurance systems: A comparison of lifetime redistribution. *Contributions to Economic Analysis*, 232:233–266, 1996.
- Marc Fleurbaey. Beyond gdp: The quest for a measure of social welfare. *Journal of Economic Literature*, pages 1029–1075, 2009.
- Marco Francesconi and Wilbert Van der Klaauw. The socioeconomic consequences of in-work benefit reform for british lone mothers. *Journal of Human Resources*, 42(1):1–31, 2007.
- Marco Francesconi, Helmut Rainer, and Wilbert Van Der Klaauw. The effects of in-work benefit reform in britain on couples: Theory and evidence*. *The Economic Journal*, 119(535):F66–F100, 2009.
- Jonathan Gruber. Disability insurance benefits and labor supply. *The Journal of Political Economy*, 108(6):1162–1183, 2000.
- Hilary W Hoynes and Erzo FP Luttmer. The insurance value of state tax-and-transfer programs. *Journal of Public Economics*, 95(11):1466–1484, 2011.

- Levell Peter Jin, Wenchao and David Phillips. A survey of the uk benefit system. 2010.
- Michael P Keane and Kenneth I Wolpin. The career decisions of young men. *Journal of political Economy*, 105(3):473–522, 1997.
- Hamish Low and Luigi Pistaferri. Disability risk, disability insurance and life cycle behavior. Technical report, National Bureau of Economic Research, 2010.
- Hamish Low, Costas Meghir, and Luigi Pistaferri. Wage risk and employment risk over the life cycle. *The American Economic Review*, 100(4):1432–67, 2010.
- Bruce D Meyer and Dan T Rosenbaum. Welfare, the earned income tax credit, and the labor supply of single mothers. *The Quarterly Journal of Economics*, 116(3):1063–1114, 2001.
- Cathal O’Donoghue. *Redistribution in the Irish tax benefit system*. PhD thesis, London School of Economics and Political Science (University of London), 2002.
- Congressional Budget Office. The distribution of household income and federal taxes, 2008 and 2009. Technical report, 2012.
- Dominique Paturot, Kirsti Mellbye, and Bert Brys. Average personal income tax rate and tax wedge progression in oecd countries. Technical report, OECD Publishing, 2013.