



## Exploring the Importance of Incentive Responses for Policy Projections

**Justin van de Ven**

National Institute of Economic and Social Research, London, UK.

Melbourne Institute of Applied Economic and Social Research, University of Melbourne, Australia.

justin.ven@simdynamics.org

**ABSTRACT:** Behavioural responses to incentives is appealing functionality for any analytical tool designed to explore the implications of public policy alternatives. This study explores the interdependence between projected savings and labour supply decisions implied by a detailed and empirically validated microsimulation model of UK households. Analysis focusses upon sensitivity of projected effects of two generic policy counterfactuals, to three alternative approaches for projecting savings and employment decisions. The results reveal that a well-specified reduced-form can generate qualitatively similar projections for policy counterfactuals to a structural approach, even if quantitative differences are difficult to avoid. Furthermore, adapting a reduced-form model to accommodate structural employment responses can be expected to obtain a close quantitative approximation to short-run projections in which both employment and savings decisions are based on utility maximisation theory. The same is not true, however, for longer-run projections due to the cumulative influence of state-specific variation.

**KEYWORDS:** DYNAMIC MICROSIMULATION, POLICY EVALUATION, SAVINGS BEHAVIOUR, LABOUR SUPPLY, STRUCTURAL, REDUCED-FORM

**JEL classification:** C51, C61, C63, H31

## 1 INTRODUCTION

Microsimulation models have been growing in scale and sophistication during recent decades, aided by improvements in data, computational capacity, and empirical and theoretical methods. One dimension of contemporary development has been the use of increasingly sophisticated methods for projecting micro-unit behaviour. In economic contexts, interest has focussed on introduction into microsimulation models of structural methods for projecting agent decisions, which provide a conceptually coherent basis for exploring the effects of policy counterfactuals. Despite improvements in the analytical tools available to model builders, however, the developmental, analytical, and computational costs of incorporating structural routines for projecting behaviour remain substantial in most realistic policy contexts. Furthermore, the advantages of including such routines are both context specific and *a priori* uncertain. In presence of clear and substantial developmental costs and opaque advantages, it is perhaps unsurprising that structural methods for projecting behaviour remain the exception rather than the rule for economic microsimulation models in use today. This paper is designed to improve the evidence base for model design, by using a detailed and empirically driven microsimulation model of UK households to explore sensitivity of projected effects of policy counterfactuals to alternative approaches for projecting labour, consumption and investment decisions.

The current treatment of behaviour in microsimulation models can be understood by putting it into historical perspective. The advent of economic modelling was made possible by two key developments during the early 1900s (Klein, 2004). First, there was the uptake of mathematical methods by economists, for both statistical evaluation (econometrics; Frisch) and theoretical development (Keynes, Hicks, Marshall, following Jevons, Menger, Walras). Secondly, the economic shocks of the great depression of the 1930s and the Second World War prompted interest in the development of public statistics.<sup>1</sup> Modern economic modelling essentially arose from a desire to make use of the newly available data via the then newly developed statistical methods, to understand the determinants of fluctuations in economic activity. Pioneering work in this regard was conducted by Jan Tinbergen, who was credited by Solow as “a major force in the transformation of economics from a discursive discipline into a model-building discipline” (Solow, 2004, p. 159).

Tinbergen’s method involved econometric estimation of a system of equations formulated to reflect theoretical insights concerning the relationship between macroeconomic variables.<sup>2</sup> Progress along this vein continued with improvements in data availability and econometric techniques. Guy Orcutt’s (1957) insight was that nonlinearities in the effects of policy on micro-units complicate projections of analyses specified at the aggregate level. He consequently advocated re-specification of economic models in terms of simulated micro-units, which could be aggregated up to macroeconomic measures. This proposition ushered in something of a golden-age for microsimulation, with intense interest in microsimulation development throughout the 1960s and 1970s.

The growing influence of econometric modelling on the policy reform process motivated associated critical appraisal of the approach. This line of enquiry culminated in growing recognition of the limitations of econometric (reduced-form) specifications for forecasting the effects of policy counterfactuals, especially from the mid-1970s (e.g. Conant & Ashby, 1970, Lucas, 1976, Campbell, 1979, and Goodhart, 1984). The source of the criticism was essentially anticipated by Keynes in his original critique of Tinbergen’s work, when he noted that “the main *prima facie* objection to the application of the [econometric] method of multiple correlation to com-

plex economic problems lies in the apparent lack of any adequate degree of uniformity in the environment” (Keynes, 1939, p. 567).

Econometric models suffered a conspicuously severe blow, when Lucas’ critique of the Philips curve for policy making purposes (an important component of econometric models of the macro-economy at the time) was generally accepted to be true (as unemployment and inflation were observed to increase in tandem).<sup>3</sup> In response to this critique, the focus of economic model development shifted from projecting decisions on the basis of econometric reduced-forms, in favour of methods based on theoretical descriptions of behaviour that are assumed to be structurally invariant to the policy environment. This shift in focus led to the development of macroeconomic models based on microeconomic foundations, with the advent of Computable General Equilibrium (CGE) models from the early 1960s, and Dynamic Stochastic General Equilibrium (DSGE) models from the 1980s.<sup>4</sup>

An outside observer might have assumed that the shift in favour of micro-foundations for economic modelling would reinforce the argument in favour of microsimulation in general. In contrast, the practical difficulties involved in using micro-foundations to project agent behaviour led modelers to adopt stylisations (e.g. representative agents) that effectively side-lined the microsimulation approach. Most developmental work on microsimulation models undertaken since the 1980s has consequently focussed on improving the statistical detail that the models capture, while retaining the traditional econometric (reduced-form) approach for projecting agent decisions. Nevertheless, there has in recent years been a noticeable trend in the microsimulation literature toward the inclusion of structural forms of decision making: In their survey of over 60 dynamic microsimulation models developed over the decade to 2013, for example, Li and O’Donoghue (2013) identify 16 that use (structural) behavioural equations to project decisions. The authors go on to note that “more models today have incorporated behavioural responses into their designs although these responses are often limited to labour market simulations” (p. 26). Hence, while models are identified that project labour supply responses to the tax-benefit system (MICROHUS, PRISM, NEDYMAS, LIAM), and others that project retirement responses to the social security system (SESIM, DYNAMITE, SADNAP), the authors conclude that there remains “limited implementation of life-cycle models in microsimulation” (p. 26).

The key difficulty with implementing a life-cycle framework within a microsimulation model is the computational burden implied by current best-practice theories. While some specifications of the life-cycle framework imply analytically convenient closed forms (e.g. Pylkkänen, 2002), these are generally ill-equipped to account for behavioural responses to uncertainty (e.g. Browning and Lusardi, 1996). Forms that do account for responses to risk generally do not have closed-form solutions, which complicates their implementation in modelling contexts.

Two alternative modelling methods have emerged that are capable of projecting behaviour where no closed form description exists. The first is based on the classical microeconomic assumption of perfect rationality, and seeks to project behaviour that optimises an assumed objective function subject to defined constraints. This approach requires computationally demanding Dynamic Programming (DP) methods to solve (e.g. Rust, 2008). In the case of DSGE models, the computational burden has been mitigated by limiting consideration to a small number of representative agents. In contrast, it was not until the 2000s that computing technology became generally available that is sufficiently powerful to permit implementation of DP methods in realistic microsimulation con-

texts. The last decade has consequently seen a growing literature based on Dynamic Stochastic Microsimulation Models (DSMMs). Much of this literature focusses on econometric evaluation of theoretical foundations, following the seminal study by Gourinchas and Parker (2002). Nevertheless, some DSMMs are starting to emerge that are designed to project the evolving national population through time for the purposes of public policy evaluation (e.g. van de Ven, 2017a).

Agent Based Models (ABMs) provide an alternative approach for projecting behaviour where no closed form description exists, by replacing the assumption of perfect rationality with a form of bounded rationality. Rather than projecting decisions that reflect a constrained optimisation, projected behaviour of each agent is based on simple heuristics or decision rules (in keeping with the keep-it-simple-stupid principle). Although originally postulated in the 1940s (von Neuman's universal constructor), ABMs did not become popular until the widespread availability of computing hardware in the 1990s. In economics, these models attracted a great deal of attention following the short-comings of CGE/DSGE models made apparent by the 2009 Global Financial Crisis (e.g. *The Economist*, July 2010). Despite growing interest, these models tend to remain highly stylised, relative to the wider microsimulation literature (see, e.g., Richiardi, 2014, and Tran, 2016; Dawid *et al*, 2016, provides an example of a more empirically orientated model in this field).<sup>5</sup>

The purpose of this paper is to clarify the implications of alternative approaches for projecting behaviour that are discussed above. Analysis focuses on the two principal behavioural margins of the domestic sector; consumption / savings, and labour / leisure. Basic economic theory suggests that savings and employment decisions are jointly determined: Stronger incentives to save can be met in part through increased labour supply; reduced returns to employment can be met in part through reduced pecuniary savings. These trade-offs are well understood and widely appreciated. Yet their implications in practical policy settings are difficult to gauge, in part because a cursory appraisal reveals them to be context specific, and in part because few models exist that permit empirical evaluation in anything approaching a realistic policy context. As a consequence, there is thin evidence for formulating adequate responses to two key modelling questions. First, generally how important is it to account explicitly for the behavioural trade-offs implied by policy counterfactuals in microsimulation projections? And secondly, to what extent will an explicit consideration of labour supply incentives alone reflect projections in which policy trade-offs concerning both the labour / leisure and consumption / savings margins are explicitly accounted for? The former of these questions addresses the overall importance of accommodating theoretical descriptions for behaviour in a microsimulation context, and the latter indicates the practical importance of analytically convenient modelling assumptions that marginalise dynamic considerations (e.g. savings) when projecting behaviour.

The current paper reports results derived from a dynamic microsimulation model that projects savings and employment decisions based on the life-cycle framework. The model is designed to project the implications of fiscal policy for the evolving population cross-section through time, and can be freely downloaded from the internet. Analysis focuses on the sensitivity of simulated effects of two policy counterfactuals, with respect to three alternative approaches for projecting savings and employment decisions. The two policy counterfactuals are a 10 percentage point increase in the rate of tax on all taxable income, and a 20 percentage point fall in the value of state retirement benefits. These counterfactuals approximate policy changes that are often considered in the literature, either because they are a focus of interest, or because they act as convenient adjustments to ensure budget neutrality of alternative reform scenarios. Savings and labour supply responses to each reform

are projected on the alternative assumptions that: i) behaviour does not respond to the influence of policy on incentives; ii) labour supply responds to policy incentives, but saving does not; and iii) both labour supply and saving respond to the incentives of policy. These alternative behavioural assumptions capture the nature of alternatives that are commonly considered in the microsimulation literature, and are explored within a single analytical framework to facilitate comparisons between them.

The remainder of the paper is comprised of three sections. The model and analytical approach are described in Section 2. Results of the analysis are reported in Section 3, and Section 4 concludes.

## 2 METHOD

The analysis is based on data generated by the Lifetime INcome Distributional Analysis model, LINDA. This microsimulation model is designed to explore the distributional implications of public policy alternatives by projecting the evolving population cross-section through time. Savings and employment are projected by the model on the assumption that these decisions are taken to maximise expected lifetime utility. This model is an ideal starting point for the current analysis, as it required only minor adjustments to generate implications of the behavioural alternatives that are the focus of interest. Version 3.16 of the model was considered for the analysis reported here, parameterised using data reported for the United Kingdom in 2011. The model is free to download from the internet at [www.simdynamics.org](http://www.simdynamics.org). This website also includes a set of video tutorials that walk-through the analysis, so that it should be possible to replicate the reported results.

An overview of LINDA is provided in Section 2.1 of the main text, and the way that the model is used to evaluate the effects of the two policy counterfactuals is described in Section 2.2. Adaptations of the model to explore alternative behavioural assumptions are discussed in Section 2.3.

### 2.1 The microsimulation model

This section provides a brief overview of the aspects of the LINDA model that were considered for analysis, and an overview of its parameterisation is provided in Appendix A. Extended discussion of the model structure is reported in van de Ven (2017b), and further detail concerning how the model parameters have been set is provided in van de Ven (2017a).

LINDA is a structural dynamic microsimulation model. It is a microsimulation model in the sense that each adult from a representative population cross-section is individually represented. It is dynamic in the sense that the model projects the characteristics of the evolving population cross-section through time. And the model is structural in the sense that labour and investment decisions are projected based on the life-cycle theory of behaviour.

LINDA considers the evolving circumstances of each adult in the evolving population cross-section, organised into annual snap-shots through time. The decision unit of the model is the nuclear family, defined as a single adult or partner couple and their dependent children. Each family is assigned a reference adult who is conceptually assumed to make all decisions on behalf of their family to maximise their expected lifetime utility, given their

prevailing circumstances, preferences, and beliefs about the future. Allocations between family members are ignored. Preferences are described by a nested Constant Elasticity of Substitution utility function. Expectations are ‘substantively-rational’ in the sense that they are either perfectly consistent with, or specified to approximate, the intertemporal processes that govern individual characteristics. The model assumes a small open economy, where rates of return to labour and capital are exogenously given (appropriate for the UK).

Utility maximising decisions were considered for consumption, labour supply, pension scheme participation, and the timing of pension access. Heterogeneity between simulated adults was limited to the following fourteen characteristics:

- |   |  |   |
|---|--|---|
| - year of birth                             | - age                                    | - relationship status <sup>u</sup>      |
| - number of dependent children <sup>u</sup> | - age of dependent children <sup>u</sup> | - student status <sup>u</sup>           |
| - education status <sup>u</sup>             | - private pension wealth <sup>d</sup>    | - timing of pension access <sup>d</sup> |
| - non-pension wealth <sup>d</sup>           | - wage potential <sup>u</sup>            | - immigration <sup>u</sup>              |
| - emigration <sup>u</sup>                   | - survival <sup>u</sup>                  |   |

Nine of the characteristics listed here are considered uncertain and uninsurable from one year to the next (represented by a  $u$  superscript), and three are projected in a way that depends on utility maximisations (represented by a  $d$  superscript).

### 2.1.1 Preferences

Expected lifetime utility of reference adult  $i$ , with birth year  $b$ , at age  $a$  is described by the time separable function:

$$U_{i,a} = \frac{1}{1-\gamma} \left\{ u \left( \frac{c_{i,a}}{\theta_{i,a}}, l_{i,a} \right)^{1-\gamma} + E_{a,b} \left[ \sum_{j=a+1}^A \beta^{j-a} \left( \phi_{j-a,a}^b u \left( \frac{c_{i,j}}{\theta_{i,j}}, l_{i,j} \right)^{1-\gamma} + (1 - \phi_{j-a,a}^b) \zeta B_{i,j}^{1-\gamma} \right) \right] \right\} \quad (1a)$$

$$u \left( \frac{c_{i,a}}{\theta_{i,a}}, l_{i,a} \right) = \left( \left( \frac{c_{i,a}}{\theta_{i,a}} \right)^{(1-1/\varepsilon)} + \alpha^{1/\varepsilon} l_{i,a}^{(1-1/\varepsilon)} \right)^{\frac{1}{1-1/\varepsilon}} \quad (1b)$$

$\phi_{j-a,a}^b$  is the probability that a reference adult with birth year  $b$  will survive to age  $j$  given survival to age  $a$ ;  $c_{i,a} \in R^+$  is discretionary composite (non-durable) consumption;  $l_{i,a} \in [0, 1]$  is the proportion of family time spent in leisure;  $\theta_{i,a} \in R^+$  is adult equivalent size based on the “revised” or “modified” OECD scale;  $B_{i,a} \in R^+$  is the legacy that reference adult from benefit unit  $i$  would leave if they died at age  $a$ ; and  $E_{a,b}$  is the expectations operator and  $A$  is the assumed maximum age that any individual can survive to. All other terms in (1) are parameters.

The assumed preference relation was selected primarily because it is standard in the associated literature.  $c$  and  $l$  are projected by the model to maximise expected lifetime utility.  $c$  is selected from a closed-and-bounded set defined to satisfy a budget constraint on liquid net wealth that is described below.  $l$  is selected from a set of

discrete alternatives, where the model was defined to permit 3 labour supply options for each simulated adult, representing full-time, part-time, and non-employment.

One innovation that has emerged in the literature, is heterogeneity of preference parameters. An early example in this vein is Gustman and Steinmeier (2005), who allow for variation in the rate of time preference to reflect heterogeneity in household retirement decisions. Variation of this type is omitted from the model to ensure that behavioural heterogeneity projected by the model is driven by heterogeneity in observable household characteristics.

### 2.1.2 Labour income dynamics

Earnings are modelled at the family level, and are described by:

$$\begin{aligned} g_{i,a} &= \max(h_{i,a}, h_{a,t}^{\min}) \lambda_{i,a} \\ \lambda_{i,a} &= \lambda_{i,a}^o \lambda_{i,a}^{emp} \lambda_{i,a}^{ret} \end{aligned} \quad (2)$$

where  $h_{i,a}$  defines the latent wage of the family of reference adult  $i$  at age  $a$ ;  $h_{a,t}^{\min}$  is the (statutory) minimum wage;  $\lambda^o$  is a random adjustment factor that is included to allow for involuntary unemployment (lack of a wage offer);  $\lambda^{emp}$  adjusts to reflect the effect of labour supply decisions on earnings (varying with endogenous  $l$ ); and  $\lambda^{ret}$  imposes a wage penalty on families that have previously chosen to start drawing upon their private pension wealth. In the analysis, the probabilities governing  $\lambda^o$  are age, relationship, and education specific, but time invariant.

In most periods, latent wages  $h$  are assumed to follow a random-walk with drift:

$$\log\left(\frac{h_{i,a}}{m_{i,a}}\right) = \log\left(\frac{h_{i,a-1}}{m_{i,a-1}}\right) + \omega_{i,a-1} \quad (3a)$$

$$m_{i,a} = m(n_{i,a}, ed_{i,a}, a, b) \quad (3b)$$

$$\omega_{i,a} \sim N(0, \sigma_{\omega}^2(n_{i,a}, ed_{i,a})) \quad (3c)$$

where the parameters  $m(\cdot)$  account for wage growth, which in turn depend on relationship status  $n_{i,a}$ , education  $ed_{i,a}$ , age  $a$ , and birth year  $b$ , and  $\omega_{i,a}$  is an identically and independently distributed disturbance term. The variance  $\sigma_{\omega}^2$  is defined as a function of relationship status and education. The only exceptions to equation (3a) are when a reference adult changes their education status (see Section 2.1.6), in which case a new random draw is taken from a log-normal distribution, the mean and variance of which are specific to the benefit unit's age, birth year, relationship, and education status.

### 2.1.3 Wealth constraint

Equation (1) is maximised, subject to an age specific credit constraint imposed on liquid (non-pension) net wealth,  $w_{i,a} \geq D_a$  for the family of reference adult  $i$  at age  $a$ .  $D_a$  is set equal to minus the discounted present value of the minimum potential future income stream, subject to the condition that all debt be repaid by age

70. Intertemporal variation of  $w_{i,a}$  is, in most periods, described by:

$$w_{i,a} = w_{i,a-1} + \tau_{i,a-1} + ur_{i,a-1}^h - c_{i,a-1} - ndc_{i,a}^x \quad (4)$$

where  $\tau$  denotes disposable income,  $ur^h$  is un-realised returns to owner-occupied housing,  $c$  is discretionary non-durable composite consumption, and  $ndc^x$  is non-discretionary expenditure. Non-discretionary costs (sometimes referred to as “committed expenditure”) are disaggregated into child care, housing (rent and mortgage interest), and ‘other’ categories to facilitate simulation of welfare benefits that make explicit reference to any one of these expenditure categories. Simulated child care costs,  $ndc^c$ , are described as a function of the number and age of dependent children, and of the employment status of the least employed adult benefit unit member. Non-discretionary housing expenditure is comprised of rent and mortgage payments,  $ndc^{hg} = rent + mort$ , and is described below. ‘Other’ non-discretionary expenditure,  $ndc^o$ , is loosely designed to reflect the minimum expenditure required to participate in society, consistent with standard definitions of poverty. Consumption on other basic necessities is defined in terms of equalised (non-housing / non-child care / non-health) consumption, and varies by age and year.

The only potential departures from equation (4) occur when a family is identified as accessing pension wealth, or when a reference adult is identified as getting married or incurring a marital dissolution. Wealth effects at the time of pension access are discussed in Section 2.1.5. In relation to marital transitions, spouses are identified from within the simulated sample. A marriage between two simulated singles consequently results in the liquid net wealth of each being combined in the common benefit unit. A divorce is assumed to see liquid net wealth split evenly between each divorcee, whereas widowhood sees all liquid net wealth bequeathed to the surviving spouse. Solutions to the utility maximisation problem are evaluated on the assumption that marriages are between identical clones.

$w$  includes all assets other than private pensions, and is disaggregated into housing and mortgage, and other wealth on the basis of reduced-form equations. Logit regressions are used to distinguish the incidence of home owners ( $hh$ ) and mortgage holders ( $mh$ ). Given incidence, regression equations for portfolio shares are used to evaluate housing wealth ( $w^h$ ), mortgage debt ( $md^h$ ), and non-housing net wealth ( $w^{nh} = w - w^h + md^h$ ). Assumed rates of return then permit evaluation of associated financial flows (realised and unrealised returns to housing wealth, mortgage interest, non-housing liquid net wealth, and rent).

#### 2.1.4 Disposable income

The model allows the measures of income accruing to each adult family member to be accounted for separately, so that it can reflect the taxation of individual incomes that is applied in the UK. Details of the specific tax and benefits schemes reflected by the model are provided in Appendix A.1. The tax function assumed for the model is represented by:

$$\tau_{i,a} = \tau \left( \begin{array}{l} b, a, n_{i,a}, n_{i,a}^c, l_{i,a}^j, g_{i,a}^j, hh_{i,a}, mh_{i,a}, w_{i,a}^h, rent_{i,a}, \\ mort_{i,a}, rr_{i,a}^h, w_{i,a}^{nh,j}, r_{i,a}^{nh}, w_{i,a}^{nh,j}, pc_{i,a}^{c/nc,j}, py_{i,a}^j, ndc_{i,a}^c \end{array} \right) \quad (5)$$



which depends on the birth year of the reference adult  $b$ ; age of the reference adult,  $a$ ; number of adults (relationship status),  $n_{i,a}$ ; number and age of all dependent children, represented by the vector  $n_{i,a}^c$ ; labour supply of each adult  $j$  in the benefit unit,  $l_{i,a}^j$ ; the labour income of each adult,  $g_{i,a}^j$ ; indicator variables for homeowners,  $hh_{i,a}$ , and mortgage holders,  $mh_{i,a}$ ; the net owner-occupied housing wealth held by the benefit unit,  $w_{i,a}^h$ ; the rent paid by non-home-owners,  $rent_{i,a}$ ; the mortgage interest paid by mortgage holders,  $mort_{i,a}$ ; the realised returns to (gross) housing wealth,  $rr^h$ ; the non-housing net liquid wealth held by each adult,  $w_{i,a}^{nh,j}$ ; the investment return on liquid net wealth of each adult,  $r_{i,a}^{nh} w_{i,a}^{nh,j}$  (which may be negative); the pension contributions made by each adult,  $pc_{i,a}^{c/nc,j}$ ; the (retirement) pension income received by each adult,  $py_{i,a}^j$ ; and non-discretionary child care costs,  $ndc_{i,a}^c$ .

### 2.1.5 Private pensions

Private pensions are modelled at the family level, and are Defined Contribution in the sense that every family is assigned an account into which their respective pension contributions are (notionally) deposited. Contributions to private pensions are defined as fixed rates of employment income conditional on (endogenous) participation, and are distinguished by whether they are made by the employer,  $\pi_{er}$ , or the employee,  $\pi_{ee}$ :  $pc_{i,a} = (\pi_{ee} + \pi_{er}) g_{i,a}$ . All employer pension contributions are assumed to be exempt from taxation, and labour income is reported net of these. Employee contributions up to a year-specific cap are also exempt from income tax, reflecting provisions of the UK tax system. Any employee contributions in excess of the cap are subject to income tax. Labour income is reported gross of all employee contributions. A cap is also imposed on the maximum size of the aggregate pension pot, which remains fixed through time.

Until the year in which a benefit unit accesses its pension wealth, intertemporal accrual of private pension wealth,  $w^p$ , is described by equation (6):

$$w_{i,a}^p = \max \left\{ 0, \min \left[ w^{p,\max}, r_{t-1}^p w_{i,a-1}^p + pc_{i,a}^p \right] \right\} \quad (6)$$

where  $w^{p,\max}$  defines the maximum size of a pension pot. Equation (6) holds in all periods prior to pension receipt except following relationship transitions, in which case associated fluctuations in pension rights are modelled in a similar fashion as described for liquid net wealth.

The age at which pension dispersals are accessed,  $a^p$ , is determined endogenously subject to a minimum age of 55 (consistent with UK policy). At the time that pension wealth is accessed, a fixed fraction of accrued pension wealth is received as a tax-free lump-sum cash payment, and the remainder converted into a level life annuity that is subject to income tax. Annuity rates are calculated to reflect birth cohort-specific survival probabilities in the model, subject to assumed rates of investment returns, real growth, and transaction costs levied at time of purchase.

### 2.1.6 Education

Each adult is allocated an education state at entry into the model,  $ed_{i,a}$ , distinguishing between those with and without graduate level qualifications. This state influences the likelihood of employment offers, the age specific evolution of latent wages ( $h$  in Section 2.1.2), and transition probabilities governing marriage and divorce.

Individuals who do not enter the simulated population with tertiary education may be identified as tertiary students,  $stud_{i,a}$ . Any individual who first appears as a tertiary student is assumed to leave tertiary education at an exogenously defined age (assuming that they survive), at which time they may transition to tertiary educated, depending on a stochastic process that represents whether they pass their final exams. At the time an individual leaves tertiary education, they receive a new random draw for their wage potential from a log-normal distribution, where the terms of the distribution differ for graduates and non-graduates. All processes that govern transitions between alternative education states when projecting a population through time are assumed to be fully consistent with the associated expectations adopted to solve the lifetime decision problem.

### 2.1.7 *Mortality*

Mortality is simulated for each adult in the model, based on random draws that are compared against associated survival probabilities. Survival probabilities are assumed to vary by age and year.

### 2.1.8 *Relationship status, spouse matching, and identification of reference adults*

A 'relationship' is defined as a cohabiting partnership, and reflects formal marriages and civil partnerships. The relationship status of each adult in each prospective year is considered to be uncertain. The transition probabilities that govern relationship formation and dissolution depend upon each reference adult's existing relationship status, their education, age, and birth year, and the mortality probability of their spouse (if one exists). These probabilities are stored in a series of 'transition matrices', each cell of which refers to a discrete relationship/education/age/birth year combination.

Relationship formations are assumed to be between members of the simulated population. At the start of each simulated year, the pool of marrying adults is identified, and sorted into couples by minimising the sum of a score that allocates one point for each year difference between simulated individuals in age, and five points for any difference in education levels. After a couple are identified, the reference adult is selected by first checking whether one partner has accessed their pension wealth but the other has not (see description of Private pensions above). If so, then the pension recipient is identified as the reference. Otherwise, the individual with the highest wage potential (see description of Labour income dynamics) is identified as the reference person.

### 2.1.9 *Children*

The model takes explicit account of the number and age of dependent children of each family. The birth of dependent children is assumed to be uncertain in the model, and described by transition probabilities that vary by the age, birth year, relationship status, and the number of existing children of each reference adult. These transition probabilities are stored in a series of transition matrices, in common with the approach used to model relationship status (described above). Having been born into a benefit unit, children are assumed to remain dependants until age 17, after which they are assumed to exit into adulthood and form family units of their own. A child may, however, depart a family prior to maturity, in the case of parental divorce. In this case all dependent children in the family are divided evenly between the separating parents (to the nearest integer).

The model is made computationally feasible by limiting child birth to three ‘child birth’ ages. Realistic benefit unit sizes are accommodated by allowing up to two children to be born at each child birth age. Restricting the number of ages at which a child can be born in the model raises a thorny problem regarding identification of the transition probabilities that are used to describe fertility risks. The model calculates the required probabilities internally, based upon the assumed birth ages and fertility rates reported at a highly disaggregated level.

#### 2.1.10 *International migration*

The model parameters include the total numbers of immigrants and emigrants to be assumed for each prospective year. The parameters also include the proportions of immigrants and emigrants to assume within a set of mutually exclusive and exhaustive population subgroups. These subgroups are defined with respect to age, education, marital status, and dependent children. Subgroups are further distinguished by disposable income quintiles for immigrants, and by past migrant status for emigrants. These model parameters permit evaluation of target numbers of immigrants and emigrants who fall into each considered population subgroup in each simulated year. The model divides the domestic population simulated for each year into the same subgroups distinguished for migrants, and randomly selects members from these subgroups as either emigrants, or to be cloned as new immigrants, to match migrant targets. Variables are generated that report the age of immigration,  $a^{im}$ , and emigration,  $a^{em}$ , for each simulated adult.

## 2.2 Evaluating the effects of policy counterfactuals

Exploring the effects of a policy counterfactual is one of the principal motivations for building microsimulation models in economics. As such, the steps involved in undertaking an associated analysis are well understood and widely appreciated. This section consequently provides a brief overview of the approach used to project the effects of each policy counterfactual. Evaluation of uncertainty associated with simulated effects of policy are then discussed, as these have received relatively sparse attention in the associated literature.

Two counterfactuals are considered for analysis. The first is an increase in the marginal rates of tax on all taxable income of 10 percentage points, which is assumed to apply in all years from 2016. Income taxes in the UK are levied on an individual basis and take a standard step-wise progressive rate structure. All income below a minimum threshold is tax exempt; see Appendix A.1.1 for further detail.

The second counterfactual is a reduction in the value of state retirement benefits of 20 percentage points, which is also assumed to apply in all years from 2016. State retirement benefits are paid from state pension age, which is scheduled to be 65 in 2016, rising to 66 from 2019, then 67 from 2026, before stabilising at 68 from 2034. This rise in state pension age is assumed for all policy environments considered in the study. State retirement benefits are defined here as the basic State Pension (a contributory flat-rate benefit; Appendix A.1.13), the Pension Credit (a means-tested retirement benefit; Appendix A.1.6), and the personal allowances for Housing and Council Tax benefits (used to evaluate means-testing of each benefit; Appendix A.1.10-11). Both counterfactuals are assumed to be announced in 2016, and to be previously unanticipated by the population.

The effects of each policy counterfactual were evaluated by comparing projections derived under the counterfactual policy environment with those derived under a base policy environment. The first step involved setting

up a simulated base scenario. In the current context, the base scenario was constructed by loading in survey data for a population cross-section reported for the UK in 2011, and projecting the evolving population cross-section forward to 2016 on the basis of a description of tax and transfer policy observed in 2011. With the population cross-section updated to 2016, LINDA was then used to project the evolving population cross-section forward from 2016 to 2070 using a description for tax and transfer policy observed in 2016 (the most recent description available at the time of writing). These simulated panel data were stored as the base for comparison.

With the base scenario in place, the effects of each of the two policy counterfactuals were evaluated by first loading in the simulated data projected under the base policy specification. These initialising data include all the information used to project individual characteristics, including the random events that evolve with uncertainty, from 2016 to 2070 under the base policy scenario. All simulated data – other than the information describing the incidence of random events – projected for the period 2017 to 2070 under the base scenario were re-initialised by the model, and new characteristics projected forward on the assumption of the respective policy counterfactual. This approach is designed to facilitate identification of the effects of policy, by limiting differences between a simulated counterfactual and the respective base projection to the policy changes of interest.

When considering the potential forecasting error associated with model projections, it is useful to distinguish between endogenous and exogenous factors that generate disparities between a model projection and the associated real-world phenomena. Endogenous factors refer to those that are explicitly recognised as evolving randomly within a model's structure. In the current model, these factors are limited to the seven individual specific characteristics that are defined as evolving with uncertainty from one year to the next (see top of Section 2.1). LINDA uses Monte Carlo methods based on random draws to project these characteristics. This means that any single projection generated by LINDA will be probabilistic (rather than deterministic) to the extent that it is affected by the specific set of random draws upon which it is based.

'Exogenous factors' refer to all other considerations that generate disparities between model projections and the associated real-world phenomena. Exogenous factors consequently capture a wide range of issues, from the uncertainty associated with defining representative model parameters, through to features that may influence the simulated phenomena but are omitted from a model's structure.

LINDA is currently adapted to account only for endogenous factors when exploring the likelihood associated with alternative model projections. This limitation makes the model an inappropriate basis for formulating a forecast concerning the future state of the population cross-section, because it means that it cannot provide a reliable measure for the full scale of the disparities associated with its projections (as it omits an account of the 'exogenous factors' referred to above).

Focussing on the effects of policy counterfactuals is motivated in part by the view that this helps to mitigate the limitations associated with use of LINDA as a forecasting tool. Comparing simulated projections that differ only with respect to the assumed policy environment, obtains a measure of the effects of policy within a controlled context. To the extent that the effects of exogenous factors are orthogonal to the effects of a policy counterfactual, taking differences between simulated projections will generate an un-biased point forecast for the effects of the counterfactual. This point forecast can then be supplemented by measures of uncertainty implied by the endogenous factors represented within the model. Conceptually, this approach is similar to the Difference-in-Differences reduced-form econometric procedure that is commonly used to explore the effects of

natural policy experiments (e.g. Blundell and Costas Dias, 2009, for a review).

LINDA is parameterised so that each simulated individual represents 1000 individuals in the projected population cross-section. Sensitivity of the projected effects of policy counterfactuals arising due to the discrete size of the simulated population is reflected by replicating the analysis 30 times, using a fresh set of random draws on each occasion. The random draws for each simulated individual are considered independent of all other individuals, so this is equivalent to bootstrapping with replacement. These replications permit standard errors to be evaluated for simulated summary statistics, which help to identify statistically significant variation associated with alternative behavioural assumptions under the base model parameterisation. Nevertheless, it should be noted that, even if the point estimates for policy effects generated as described here are un-biased, the estimated uncertainty associated with those point estimates is likely to be substantively under-predicted due to the omission of the uncertainty of model parameters in the associated calculations.

### 2.3 Exploring alternative behavioural assumptions

One of the most flexible methods for projecting decisions is on the basis of reduced-form descriptions of behaviour. In its most general form, the approach involves the assumption of a functional relation between the behaviour of interest and alternative 'explanatory' characteristics. The central problem with this most general form is the difficulty associated with identifying a functional relation that is stable in context of an evolving policy environment. A model is appropriately referred to as 'structural' to the extent that this objective is achieved. Hence, a single model can be referred to as both 'structural' and a 'reduced-form'. Unfortunately, not all reduced-form descriptions are structurally invariant to the policy environment, and not all structural descriptions of behaviour can be expressed as convenient reduced-forms.

Choice of the method used to project behaviour is currently part of the art – rather than the science – of specifying a microsimulation model. The choices made in this respect tailor a model to selected subjects of interest. The multiplicity and indeterminacy of this aspect of model design also limits the scope of comparisons that can reasonably be explored between modelling alternatives; it is impractical, for example, to consider the question of which approach is likely to be 'best' for all contexts. The subject explored in this study was carefully chosen to be both informative and feasible.

One of the most difficult features to reflect when selecting a reduced-form description for behaviour is incentive responses to a changing policy environment. Understanding such responses is the primary motivation for utility theory, and a central focus of concern in the growing field of agent-based economics. This study consequently uses a microsimulation model based on classical utility theory to explore the importance of incentive responses underlying projections for policy counterfactuals. The analysis explores the influence of three alternative behavioural scenarios: (a) Behaviour that does not respond to the influence of policy on incentives; (b) labour supply that responds to altered policy incentives but savings behaviour that does not; and (c) the case in which both savings and labour supply respond to changes in policy incentives. These behavioural alternatives are interesting because they reflect assumptions frequently made in the contemporary microsimulation literature.

Use of a common model framework facilitates comparisons between the three alternative behavioural assump-

tions considered for the study. The framework adopted for the analysis, LINDA, has been designed and parameterised exclusively on the assumption that both savings and labour supply respond optimally to changes in the policy environment, reflecting one of the three behavioural alternatives listed above (c). This tailoring of the considered model framework motivates use of the “fully-structural” projections (c) as the basis for comparing the behavioural alternatives that reflect reduced-forms (a,b). The analysis is consequently framed around two principal research questions: How important is it to account explicitly for the behavioural trade-offs implied by policy counterfactuals in microsimulation projections (a,c); and to what extent will an explicit consideration of labour supply incentives alone reflect projections in which policy trade-offs concerning both the labour / leisure and consumption / savings margins are explicitly accounted for (b,c)?<sup>6</sup>

The remainder of this section provides technical detail concerning the methods used to project behavioural alternatives.

### *2.3.1 Structural projections for behaviour using LINDA*

LINDA projects the evolving population cross-section through time via a two stage process that is common in the dynamic programming literature. In the first stage, utility maximising decisions are evaluated under the simulated policy environment for any feasible combination of individual specific characteristics. In the second stage, the population is projected through time, based on the description of behaviour evaluated in the first stage, and the processes that are assumed to govern the intertemporal development of individual specific characteristics. The second stage of this procedure is common to most dynamic microsimulation models in use today, with the principal distinction that LINDA uses the utility maximising solutions obtained in the first stage for projecting individual decisions. Exploring the implications of alternative behavioural assumptions required adaptation of the description for behaviour generated in the first stage of the model procedure, which is the focus of discussion here.

The model begins by defining a grid that overlays all feasible combinations of individual specific characteristics. This grid essentially acts as a look-up table for projecting individual decisions in the second stage of the simulation. A single grid axis is generated for each of the 14 characteristics described at the beginning of Section 2.1. Where a characteristic defines discrete alternatives (e.g. relationship status, number and age of dependent children, student and education status, and survival), then a separate point is defined on the respective grid axis for each potential alternative. Where a characteristic defines a continuous feasible domain, then the respective axis is arbitrarily divided into discrete points. For characteristics concerning time, a linear scale with intervals of equal duration is adopted. Annual intervals are considered for the age axis, comprising 113 points between 18 and 130 years. The working lifetime is considered to end at age 74, and all individuals are assumed to be retired from age 75. Feasible birth years span the period 1920 to 2030, and this period is divided into 11 year intervals (11 points). For financial characteristics (wage potential, pension and non-pension wealth), equal intervals on a log scale are used to construct the grid, as this provides greater detail toward the bottom of the distribution where non-concavities of the utility maximisation problem are most common. The feasible domains for non-pension wealth and wage potential are each divided into 26 points to age 74, and 21 points for pension wealth. From age 75, 151 points are considered for each of non-pension and pension wealth where greater detail can be accommodated at a low computational cost. In total, the grid considered for analysis identifies 200 million discrete

combinations of individual specific characteristics over the feasible range of the analysis.

Once the grid described above has been defined, the model proceeds to solve the utility maximisation problem at each grid intersection via backward induction. Starting with the oldest possible age considered for life (130 years), utility maximising decisions for any combination of individual specific characteristics (birth cohort, pension and non-pension wealth) are trivial to evaluate, as they are free from dynamic considerations. The model stores both the utility maximising decisions and the respective measures of (lifetime) utility at each intersection of the grid slice for the oldest age, and then proceeds to consider the decision problem at each intersection of the grid slice for the second oldest potential age of life. Solution of the utility maximisation problem at age 129 is complicated by the need to take into consideration the impact that decisions at age 129 have on probable characteristics, and associated utility, at age 130.

As closed-form (analytical) solutions to the intertemporal utility maximisation problem are not guaranteed in the model, numerical methods are used to search over the set of all feasible decisions. Expected lifetime utility associated with each candidate decision combination is evaluated by: i) calculating the period specific utility associated with the decision, via equation (1b); ii) projecting characteristics forward one year taking into consideration the influence of the decision on future circumstances, e.g. via equation (4) for liquid net wealth; and iii) by approximating the expected lifetime utility associated with each potential forward projection for individual specific characteristics by drawing on the solutions previously obtained for the grid slice describing the immediately succeeding age. Where forward projections of characteristics are uncertain, then a discrete set of alternatives are considered, each associated with a probability. In the case of wages, where the theoretical distribution one year forward is continuous, an approximation is obtained using the Gauss-Hermite quadrature. Where one year forward projections for characteristics do not fall precisely on a grid point, then interpolation methods are used to approximate the measure of expected lifetime utility one-period forward by drawing on near-by grid points. The assumption of von Neuman Morgenstern utility then permits expected lifetime utility associated with the decision combination to be evaluated as a weighted sum.

As the model solves each utility maximisation problem, it stores both the optimal decision combination and associated measure of expected lifetime utility. Once it has a solution for all intersections of the grid slice for age 129, it uses the same approach applied recursively to obtain solutions for the entire simulated lifetime. At the end of this process, the model stores the results of its solution to the lifetime decision problem to an external file, which plays an important role in facilitating the reduced-form analysis as discussed below.

### 2.3.2 *Adapting LINDA to explore reduced-form behavioural alternatives*

All three behavioural variants considered in the study explore the effects of policy counterfactuals with respect to a common simulation base (see Section 2.2). Given the focus on behavioural reduced-forms that do not respond to the incentives associated with a change in policy context, the reduced-form projections use descriptions for behaviour described by the base simulation. These descriptions for behaviour were generated by LINDA via the structural approach described in Section 2.3.1, taking into consideration transfer policy as it was applied in the UK in 2016.

The descriptions of behaviour generated by LINDA under the base policy environment take the form of a

(look-up table) function, with exogenous variables defined by the 14 household-specific characteristics listed at the beginning of Section 2.1. A common feature of the 14 exogenous characteristics assumed for the model is that their values in any simulated period are independent of the policy environment prevailing in the same period. This feature is useful, because it means that the characteristics generated under a base simulation can be used as the starting-point for projections of a policy counterfactual. In this context, immediate responses to the policy counterfactual are accommodated in numerically solved structural projections by resolving the (lifetime) decision problem for each alternative policy environment. In contrast, reduced-form descriptions of behaviour are usually assumed to persist in context of a changing policy environment. This difference in simulation approaches complicates a direct interchange of the respective descriptions for behaviour.

Consider, for example, projections of consumption decisions in context of a rise in income tax rates. Suppose that consumption was projected using a functional description based on age, wealth, and private income. The structural approach that is implemented by default in LINDA would adapt the consumption function to reflect the influence on incentives of higher tax rates. This approach consequently accommodates an immediate reaction to the policy change. In contrast, if the reduced-form consumption function was not revised to reflect the income tax rise, and private income remained unchanged, then projected consumption would fail to reflect the higher taxes until budgetary effects fed through to affect household wealth. The result is likely to be a run-away wealth effect, and projections for consumption that bear little resemblance to reality.

Studies that employ reduced-form descriptions of behaviour must carefully tailor their functional descriptions to minimise the likelihood of mis-specification as discussed above. This is typically achieved, within the constraints considered for reduced-forms here, by reflecting relationships between behaviour and explanatory variables that are likely to suffer little distortion by the policy environment. In the above example, this would typically be achieved by describing consumption as a function of disposable, rather than private income. A rise in tax rates would then depress disposable income for any given measure of private income, and thereby have an immediate effect on consumption, independent of any associated implications for wealth. While the magnitude of consumption responses generated under the reduced-form and structural approaches might not be the same, this approach for selecting reduced-form specifications is likely to mitigate the differences between them.

The approach assumed here for projecting behaviour under the reduced-form was consequently chosen to reflect a conceivably stable description in context of the two policy counterfactuals that are explored. As discussed above, this reflects basic best-practice in reduced-form analysis. Much of the indeterminacy commonly associated with specification of a reduced-form description was addressed here by imposing the additional condition that the reduced-form analysis retained the basic “information set” underlying the structural description for behaviour. This additional condition was imposed to facilitate comparisons between the two methodologies, and involved adapting the axes of the look-up tables to reflect the bearing of each policy counterfactual on disposable, rather than private income. This was achieved as follows.

First, copies were made of the look-up tables describing the behavioural solutions generated by LINDA under the base policy environment. Each coordinate  $i$  of the base-simulation look-up table was considered in turn, and the ‘wage potential’ and ‘private pension wealth’ axes were considered for adjustment. These two axes were adjusted independently, reflecting the fact that households predominantly receive income from either one, but seldom both, of these income sources in any given year of their lives.



Re-scaling of the ‘wage potential’ axis was evaluated for any coordinate,  $i$ , corresponding to the working lifetime (to age 75). Given the characteristics associated with coordinate  $i$ , including wage potential  $x_i$ , disposable income under the policy counterfactual,  $y_i^1$ , was evaluated, on the assumptions that all adults were full-time employed and no private pension was received. The ‘wage potential’ needed to obtain the same disposable income, under the same assumptions, except in context of the base policy context,  $x_i^0$ , was then evaluated. This was used as the reference for re-scaling the ‘wage potential’ of the respective coordinate in the look-up table used to project reduced-form behaviour. That is, the reduced-form projections for behaviour of individuals with wage potential  $x_i$  under the policy counterfactual were drawn from the base-simulation look-up tables, with reference to wage potential  $x_i^0$  rather than  $x_i$ .

Consider, for example, the re-scaling associated with the counterfactual 10 percentage point increase in all income tax rates. Denote  $t_i$  the marginal tax rate under the base simulation corresponding to wage potential  $x_i$ ,  $\Delta t = 0.1$  the rise in tax rates, and the proportion of wage potential subject to tax  $\phi_i = (x_i - e_i)/x_i$  (where  $e_i$  is exempt income). Then it can be shown that the re-scaled value of the wage potential at coordinate  $i$  is given by:<sup>7</sup>

$$x_i^0 = \left(1 - \frac{\Delta t}{(1 - t)} \phi_i\right) x_i \quad (7)$$

The approach used to re-scale the ‘private pension wealth’ axis adapted the method described above to capture the forward-looking nature of pension savings decisions. Any coordinate identified as having previously accessed pension wealth provides a direct measure of the associated (private) pension annuity. For coordinates identified as not having previously accessed pension wealth, a private pension annuity was imputed. This imputation was evaluated either at state pension age, or the age of the table coordinate, whichever was higher. The imputation accounted for projected investment returns, the (exogenous) share of pension wealth taken as a tax-free lump-sum, and annuity conversion rates. Given the private pension annuity, disposable income under the policy counterfactual,  $y_i^1$ , was evaluated. The private pension annuity needed to obtain the same disposable income, under the same assumptions, except in context of the base policy context,  $x_i^0$ , was then determined.<sup>8</sup> The analytical process was then inverted to obtain a measure of private pension wealth at the prevailing coordinate, and used as the reference to re-scale the ‘private pension wealth’ axis in the look-up table adapted from the base simulation.

The adjusted look-up tables described above were used to generate counterfactual policy projections under the behavioural alternative in which both savings and employment decisions are based on reduced-forms. Consideration of the intermediate behavioural scenario, in which labour supply adjusts to the altered incentives associated with a policy counterfactual but savings behaviour does not, required some additional adjustment of the existing LINDA source code. In this case, the model was adapted to first load in the adjusted look-up tables described above. The numerical routines that search for an optimum to the utility maximisation problem were altered to project consumption and pension savings using the pre-loaded values, and to search only over the feasible labour supply alternatives. When the search considered the alternative of no labour supply, then consumption and pension participation decisions were set equal to the values associated with a low-wage offer (see Section 2.1.2); otherwise consumption and pension participation were set equal to their values with a standard wage offer. The results of this constrained optimisation were saved by the model, after which each simulation

proceeded as normal.

Results based on the adjusted look-up tables obtained as described above are reported in Section 3. Results for reduced-form projections based on unadjusted look-up tables are reported in Appendix B, thereby providing an indication of the sensitivity of results to the re-scaling adjustments considered for analysis; these supplementary results clearly display the forms of mis-specification that are discussed above.

### 3 RESULTS

Results are divided into two subsections. The first provides an overview of the influence of alternative behavioural assumptions, by discussing the effects of the respective policy counterfactuals on key aggregates generated for the evolving population cross-section. The second subsection explores the dynamics underlying the population aggregate effects, by discussing the influence of alternative behavioural assumptions on the projected life course of selected birth cohorts.

#### 3.1 Policy effects on the evolving population cross-section

Table 1 reports the projected effects on the government budget of a 10 percentage point rise in all rates of income tax, and a 20 per cent fall in the value of state retirement benefits. As discussed in Section 2.3, the absolute scale of differences between behavioural alternatives reported for each policy counterfactual reflect this study's approximation of the influence of accommodating incentive responses to policy based on a common information set. A cursory examination of the statistics reported in the table indicates a reasonably high degree of correspondence between statistics generated under the three behavioural alternatives, with a closer correspondence projected for the fall in retirement benefits than the rise in tax rates. This disparity in behavioural variation is generated despite the two policy counterfactuals being projected under the same set of behavioural alternatives, and having qualitatively similar effects on net government revenues. The result is an important reminder of the extent to which behavioural sensitivity of simulated projections is policy-specific.

The top-left panel of Table 1 indicates that the rise in tax rates is projected, under the assumption of no responses to counterfactual incentives (the 'non-response' scenario, denoted 'none' in the table), to increase net government revenue by £34 billion in 2016, rising to £37 billion by 2046. The figure reported for 2016 is substantively lower than related projections reported by HM Revenue and Customs (HMRC), which suggest that a 1 percentage point rise of all rates of income tax in the 2016/17 tax year would raise £5.5 billion in additional revenue. The reason for this disparity is that the reduced-form descriptions for behaviour assumed here are explicitly designed to reflect immediate labour supply responses throughout the distribution, in common with the 'full' structural approach for projecting behaviour (discussed below). In contrast, the HMRC projection confines responses to higher-rate tax-payers, assuming that a 1 per cent rise in the marginal tax rate would lead to a 0.45 per cent reduction in taxable income. This disparity consequently reveals the potential sensitivity of reduced-form projections based on heuristically selected descriptions for behaviour.<sup>9</sup>

The top-left panel of Table 1 also indicates that the increases in net government revenue projected for the rise in tax rates under the 'full-response' scenario (in which both savings and labour supply respond to incentives) are

**Table 1:** Projected effects of policy counterfactuals on annual government budget, distinguished by year and accommodated behavioural response (£2016, billions).

	income tax rates rise 10%			retirement benefits fall 20%			income tax rates rise 10%			retirement benefits fall 20%		
	none	emp	full	none	emp	full	none	emp	full	none	emp	full
	net government revenue						benefits expenditure					
2016	33.5	39.5	41.7	16.6	15.7	17.3	0.9	0.5	0.5	-20.1	-19.8	-19.7
	(0.4)	(0.4)	(0.3)	(0.1)	(0.1)	(0.1)	(0.1)	(0.0)	(0.0)	(0.1)	(0.1)	(0.1)
2021	33.0	40.0	44.4	19.0	18.1	19.3	1.5	0.4	0.6	-23.0	-22.6	-22.3
	(0.5)	(0.4)	(0.4)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
2026	32.4	39.9	45.1	22.4	21.3	22.1	1.5	0.3	0.7	-26.4	-26.0	-25.5
	(0.4)	(0.4)	(0.4)	(0.2)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
2036	31.0	39.6	47.1	33.7	31.4	31.8	2.0	0.0	0.9	-37.3	-36.9	-36.0
	(0.4)	(0.4)	(0.5)	(0.2)	(0.2)	(0.2)	(0.2)	(0.1)	(0.1)	(0.2)	(0.2)	(0.2)
2046	36.6	44.0	54.9	46.8	42.8	43.9	2.1	-0.3	1.1	-50.2	-49.7	-48.6
	(0.5)	(0.5)	(0.5)	(0.2)	(0.2)	(0.2)	(0.1)	(0.1)	(0.1)	(0.2)	(0.2)	(0.2)
	income tax revenue						consumption tax revenue					
2016	40.9	46.3	44.9	-0.1	-0.9	-1.3	-6.5	-6.3	-2.7	-3.4	-3.1	-1.1
	(0.4)	(0.4)	(0.3)	(0.1)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
2021	41.2	46.7	48.7	-1.0	-1.6	-1.8	-6.7	-6.3	-3.6	-3.1	-2.9	-1.1
	(0.5)	(0.4)	(0.5)	(0.1)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
2026	40.9	46.6	50.1	-1.2	-2.0	-2.2	-7.0	-6.4	-4.2	-2.8	-2.7	-1.1
	(0.4)	(0.4)	(0.4)	(0.2)	(0.1)	(0.1)	(0.0)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)
2036	41.0	46.9	53.3	-1.1	-2.8	-2.9	-8.0	-7.3	-5.3	-2.6	-2.7	-1.3
	(0.4)	(0.4)	(0.5)	(0.2)	(0.1)	(0.1)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)
2046	48.8	53.0	62.6	-0.6	-3.9	-3.3	-10.0	-9.2	-6.6	-2.7	-3.0	-1.4
	(0.5)	(0.5)	(0.5)	(0.2)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)

Notes: \* “none” = projections omitting behavioural responses to policy incentives; “emp” = projections allowing for labour supply responses to policy incentives; “full” = projections allowing for labour and savings responses to policy incentives. Standard errors reported in parentheses. “income tax rates rise 10%” denotes counterfactual in which the marginal rates on all taxable income are increased by 10 percentage points. “retirement benefits fall 20%” denotes counterfactual in which all state retirement benefits are reduced in value by 20 percentage points.

Source: Author’s calculations on simulated data generated using 30 separate sets of random draws.

significantly greater than under the non-response scenario, with the excess rising from 20 per cent in 2016 to 33 per cent in 2046. The remaining three panels of the table help to explain this variation by disaggregating the net budgetary effects. These statistics indicate that all three budgetary components contribute to the differences in net government revenue projected for the rise in income tax rates under the full- and non-response scenarios. That is, the higher increases in net government revenues projected under the assumption that both savings and employment respond to policy incentives are coincident with smaller increases in benefits expenditure, larger increases in income tax revenue, and smaller decreases in consumption tax revenue. The most important of these components, by a wide margin, is income tax revenue, followed by consumption taxes; differences in projected benefits account for relatively little of the variation projected for net government revenues.

The relative importance of income taxes in contributing toward the differences in the effects of a rise in tax rates on the net government budget projected under the non- and full-response scenarios alludes to the likely importance of labour decisions as an explanatory factor.<sup>10</sup> This proposition suggests that extending the non-response scenario to accommodate labour supply responses to policy incentives may obtain a close approximation to the projections generated under the full-response scenario.

The statistics reported in Table 1 for simulations in which labour supply is assumed to respond to the altered incentives of the rise in income tax rates but savings do not (the ‘employment-response’ scenario, denoted ‘emp’ in the table) generally fall between those generated under the two behavioural alternatives, thereby providing some support for the conjecture noted in the previous paragraph. Although this result may seem intuitive, it is nevertheless useful to confirm, and provides some support for the prevailing trend to add structural labour supply responses to microsimulation models that include reduced-form descriptions for consumption. It is of further note that the intermediate statistics generated under the employment-response scenario reported in the table tend to be closer to those generated under the full-response scenario in the near-term (within a 10 year horizon), and closer to the non-response scenario in the longer term. It is conceivable that policy makers would find the relative variation generated under the employment-response scenario strictly preferable to either of the two behavioural alternatives reported in the table. We return to discuss the generality of this finding below.

While the effects of alternative behavioural assumptions on projected net government revenues are relatively muted under the retirement benefits counterfactual, some interesting differences between the projections are discernable. The increase in net government revenues projected for 2016 is £700 million larger under the full-response scenario than under the non-response scenario. This relationship reverses as the simulated time horizon extends, so that by 2046 the projected increase in net government revenue is £2.9 billion smaller under the full-response scenario. The remaining statistics in Table 1 indicate that £1.2 billion of this temporal reversal is due to a smaller decline in state expenditure on retirement benefits under the full-response scenario, £1.4 billion is due to a larger fall of income tax revenue, and £1.0 billion is due to a larger decline in consumption tax revenue. Further detail can be obtained by exploring the contemporaneous variation generated for household finances, to which we now turn.

Table 2 reports policy effects on domestic sector financial aggregates. Most of the statistics reported in the table indicate similar directional effects of each reform across the three behavioural alternatives, subject to some notable differences in magnitude.

Starting with statistics reported for the effects of the rise in income tax rates, labour supply is projected to fall

**Table 2:** Projected effects of policy counterfactuals on aggregate domestic sector finances, distinguished by year and accommodated behavioural response (£2016 billion, unless otherwise stated).

	income tax rates rise 10%			retirement benefits fall 20%			income tax rates rise 10%			retirement benefits fall 20%		
	none	emp	response* full	none	emp	full	none	emp	response* full	none	emp	full
	equivalent full time employees ('000)						consumption expenditure					
2016	-1481.6 (21.1)	-491.8 (7.4)	-574.9 (9.8)	307.8 (10.7)	237.9 (7.6)	160.8 (5.6)	-43.0 (0.3)	-41.8 (0.3)	-18.7 (0.2)	-22.4 (0.1)	-20.6 (0.1)	-7.5 (0.0)
2021	-1300.9 (18.5)	-459.6 (11.4)	-345.1 (12.3)	284.4 (13.1)	194.7 (10.3)	148.5 (8.4)	-44.6 (0.3)	-42.0 (0.3)	-24.9 (0.2)	-20.8 (0.1)	-19.7 (0.1)	-7.7 (0.1)
2026	-1198.7 (19.4)	-435.1 (16.0)	-218.9 (14.4)	321.6 (14.9)	183.0 (10.6)	153.0 (9.7)	-46.9 (0.3)	-43.1 (0.4)	-29.0 (0.3)	-19.1 (0.1)	-18.5 (0.1)	-8.0 (0.1)
2036	-1186.1 (20.1)	-500.0 (17.7)	-124.6 (17.8)	396.1 (14.8)	173.9 (13.0)	173.3 (11.4)	-54.4 (0.4)	-49.5 (0.4)	-36.1 (0.3)	-18.2 (0.2)	-19.0 (0.2)	-9.5 (0.2)
2046	-1081.9 (20.3)	-614.6 (17.5)	-156.1 (14.2)	430.4 (19.2)	105.6 (12.1)	176.9 (12.1)	-68.3 (0.5)	-62.4 (0.5)	-45.7 (0.4)	-19.2 (0.3)	-21.5 (0.3)	-10.8 (0.2)
	private income						disposable income					
2016	-33.1 (0.7)	-15.2 (0.4)	-17.4 (0.4)	7.9 (0.3)	5.3 (0.3)	3.5 (0.2)	-70.7 (0.6)	-59.1 (0.5)	-61.8 (0.5)	-13.8 (0.2)	-14.9 (0.2)	-15.8 (0.1)
2021	-34.2 (0.6)	-15.3 (0.5)	-13.1 (0.5)	8.0 (0.4)	6.1 (0.3)	3.0 (0.2)	-71.5 (0.6)	-59.4 (0.5)	-61.4 (0.5)	-15.7 (0.3)	-16.3 (0.2)	-18.3 (0.2)
2026	-37.0 (0.7)	-16.5 (0.6)	-12.6 (0.5)	11.3 (0.6)	8.1 (0.4)	3.6 (0.3)	-74.0 (0.7)	-60.7 (0.7)	-62.7 (0.7)	-16.0 (0.4)	-17.6 (0.3)	-20.8 (0.2)
2036	-47.7 (0.8)	-23.9 (0.8)	-15.1 (0.7)	18.9 (0.6)	11.3 (0.4)	4.8 (0.4)	-84.3 (0.7)	-68.4 (0.7)	-68.7 (0.7)	-20.2 (0.5)	-25.0 (0.4)	-29.7 (0.4)
2046	-61.7 (1.2)	-40.0 (0.9)	-25.0 (0.8)	27.2 (0.7)	12.3 (0.6)	6.5 (0.5)	-105.1 (1.2)	-90.0 (0.9)	-87.5 (0.8)	-25.4 (0.6)	-35.5 (0.7)	-40.5 (0.6)
	pension wealth						non-pension wealth					
2016	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
2021	-35.8 (2.1)	-36.4 (1.4)	-8.6 (1.7)	45.8 (2.4)	25.4 (1.6)	17.4 (1.4)	-102.1 (2.4)	-48.6 (1.4)	-175.9 (1.8)	41.4 (2.3)	39.2 (1.7)	-37.8 (1.5)
2026	-69.2 (2.9)	-77.0 (1.7)	-9.1 (2.4)	89.5 (3.2)	50.7 (1.6)	34.4 (1.5)	-206.7 (3.1)	-103.7 (2.3)	-326.9 (3.1)	83.5 (2.7)	72.8 (1.6)	-78.4 (1.6)
2036	-125.8 (3.7)	-142.3 (2.6)	17.6 (3.0)	201.1 (4.0)	118.8 (2.8)	78.6 (2.5)	-408.3 (6.0)	-220.3 (4.0)	-585.5 (4.7)	150.1 (4.6)	112.6 (3.8)	-170.7 (3.2)
2046	-161.9 (6.9)	-211.0 (5.8)	56.0 (5.4)	366.1 (8.0)	210.7 (3.4)	140.2 (3.0)	-646.3 (9.4)	-360.4 (6.3)	-863.8 (8.1)	236.4 (5.9)	142.8 (4.5)	-280.4 (4.5)

Notes: \* "none" = projections omitting behavioural responses to policy incentives; "emp" = projections allowing for labour supply responses to policy incentives; "full" = projections allowing for labour and savings responses to policy incentives. Standard errors reported in parentheses. "equivalent full-time employees" evaluated as projected change in total number of labour hours per week divided by 37. "private income" denotes income net of interest charges from all private sources. "disposable income" denotes income net of government taxes and transfers.

Source: Author's calculations on simulated data generated using 30 separate sets of random draws.

following the policy change under all three behavioural alternatives, with the scale of the reduced employment declining over the simulated time horizon. The largest effects (by a wide margin) are reported under the non-response scenario, where the number of (equivalent) full-time employees is projected to fall by almost 1.5 million (from a total population of 53 million adults) following the rise in income tax rates in 2016, declining to a fall of just over 1 million by 2046. In contrast, employment under the full-response scenario is projected to decline by just over half a million in 2016, and by less than 200 thousand by 2046. As with the effects reported for the government budget, the employment-response scenario lies between these two extremes, is closer to the full-response scenario early in the projected time horizon, and drifts toward the non-response scenario with time.

The rise in tax rates is associated with a price effect that discourages employment, and a wealth effect that encourages employment. The relatively large fall projected for employment under the non-response scenario can consequently be understood as a reflection of exaggerated price effects captured by the assumed reduced-form description for behaviour. As discussed in Section 2.3, the behavioural description assumed for the non-response scenario was obtained by re-scaling the look-up table generated by the structural model under the base policy environment to capture the effects of each reform on disposable, rather than private income. In terms of wage potential, the re-scaling was evaluated by assuming full employment of all adult household members. This assumption was made to reflect the impact of each counterfactual on lifetime circumstances, in a way that is independent of the actual labour supply decisions; essentially, this approach was selected to capture the wealth effect, given the assumed information set.

As the approach adopted for identifying a reduced-form behavioural description was arbitrarily defined, it is possible that an alternative identification approach could mitigate the behavioural sensitivity of projected effects described in Table 2. Results presented in Appendix B, for example, indicate that projecting reduced-form behaviour on the un-scaled look-up tables calculated under the base policy environment generates rates of labour supply following the rise in tax rates that are substantively *higher* than under the full-response scenario. Adopting a reduced-form behavioural description intermediate between these two alternatives may consequently obtain a closer match between projections for labour supply generated under the full- and non-response scenarios. Searching for a reduced-form behavioural description that obtains a closer match to the full structural projection would, however, miss an important point.

The wealth effect that dampens the fall in labour supply projected under the full-response scenario is the product of a forward-looking expectation in context of the assumed policy counterfactual. This expectation is unobserved at the time the counterfactual first takes effect, and is consequently difficult to capture in a reduced-form. Interpreted in this way, the disparities in employment effects reported for the rise in income taxes between the non- and full-response scenarios are a product of the incentive effects that are the focus of the current analysis. While approximations might be devised to reflect these incentives within a reduced-form – one example being explored below for the fall in state retirement benefits – generic formulations that are easily implemented and suitable for a wide-range of policy alternatives do not exist, and so are usually omitted from microsimulation contexts.

The projected effects of the higher tax rates on labour supply discussed above feed directly through to influence private income. Disposable income falls by a larger margin than private income, due to the co-incident rise in taxes paid. The declines projected for disposable income generate, by construction, coincident declines in con-

sumption under the non-response scenario. While the full-response scenario also projects declines in consumption under the higher tax policy, these declines are smaller than those reported for the non-response scenario. This is partly attributable to the smaller declines projected for disposable income under the full-response scenario. Importantly, however, it also reflects the reduction in savings incentives associated with the higher tax environment, which are not taken into consideration by the reduced-form analysis. These savings responses can also be seen in the wealth effects of the rise in taxes reported in Table 2, which indicate a rise in (tax-shielded) pension wealth, more than off-set by a relative decline in non-pension wealth under the full-response scenario.

The consumption effects reported for the employment-response scenario are conspicuous, in context of the preceding discussion, by the close correspondence that they display with the non-response scenario throughout the simulated time horizon. This is unsurprising, as consumption is projected on the same basis under both the non- and employment-response scenarios. Nevertheless, it is useful to note, as it is responsible for the relative magnitudes of effects generated under the three behavioural alternatives.

As discussed previously, projected effects reported for the government budget in Table 1 under the employment-response scenario tend to be intermediate between the full- and non-response scenarios, starting closer to the full-response scenario and tending toward the non-response scenario as the time horizon is extended. This is potentially appealing variation, especially for those who are sceptical of the empirical validity of theoretical descriptions of behaviour. Table 2 reveals that, in the near-term, projections for the government budget under the employment-response scenario tend to be closer to the full-response scenario because of the closer correspondence of labour supply responses. It seems reasonable to suppose that this close correspondence between the full- and employment-response scenarios will hold for a broad class of policy reforms, because the behavioural problem for both the full- and employment-response scenarios essentially share a common state-space.

However, the current analysis reveals that the longer term relationship between the three behavioural alternatives will depend upon quantitative differences between reduced-form and structural projections for labour/leisure and consumption/savings. In the current context, for example, the reduced-form projections imply a larger reduction in consumption and employment under the higher tax regime than the structural projections. Lower consumption implies higher future savings, which generally depresses employment, and it is this process that generates the drift of the employment-response scenario toward the non-response scenario over the simulated time horizon. In general, it is not usually known how these behavioural alternatives relate, so that it is impossible to say how a projection that accommodates employment responses to incentives in context of a reduced-form description for consumption/saving might vary, relative to a full structural framework. What is clear, however, is that quantitative distortions associated with a reduced-form description for consumption can generate substantial biases in longer-term employment projections, even if the structural description for employment is well-specified. This is an important observation, given the cursory consideration that is sometimes paid to reduced-form descriptions for consumption when accommodating structural labour-supply decisions.

In contrast to the counterfactual rise in income tax rates, the fall in retirement benefits is associated with price and wealth effects that work in the same direction, and price effects that are relatively weak. The method used to identify reduced-form behavioural responses to the fall in state retirement benefits is consequently adapted to capture the influence of the reform on expected returns to pension saving, as discussed in Section 2.3. Table 2 suggests that the identification strategy was successful, in the sense that broadly similar effects were projected for

the policy counterfactual across all three behavioural alternatives. All three behavioural scenarios project falls in leisure and consumption throughout the simulated time horizon, as households work and save more to off-set the reduced generosity of state benefits. The associated increases in employment generate higher measures of private income, which are more than off-set by coincident declines in state retirement benefits paid, so that disposable income for population cross-sections is projected to fall in all simulated years and behavioural scenarios. The increased saving among the working aged population appears as a rise in pension wealth in all projections. The only directional differences between behavioural scenarios are reported for non-pension wealth, which declines under the full-response scenario, but increases under the reduced-form alternatives.

Relative to the non-response scenario, the full-response scenario projects a smaller increase in labour supply, and a smaller decline in consumption following the reduction in state retirement benefits throughout the simulated time horizon. These differences in behavioural response generate lower projections for wealth under the full-response scenario, relative to the non-response scenario. The implication is that the approach used to account for the wealth effect of the counterfactual reduction in state retirement benefits when identifying the reduced-form behavioural descriptions exaggerates the effect described by the full-structural projection. This highlights the difficulties associated with heuristically identifying a reduced-form specification.

In common with the projections for the rise in tax rates, projections under the employment-response scenario for consumption responses to the fall in retirement benefits are similar to those projected under the non-response scenario, and the employment responses generally fall between the two behavioural alternatives. In contrast to the counterfactual tax analysis, however, the labour effects projected under the employment-response scenario diverge from the non-response scenario with the simulated time horizon. The reason for this is that the reduced-form consumption responses to the fall in state retirement benefits tend to exaggerate savings, which depresses labour supply, while the reduced-form employment responses tend to exceed those generated by the structural model. This policy counterfactual consequently provides a nice example of the uncertainty associated with temporal biases of employment projections based on reduced-form descriptions for consumption as discussed above.

### **3.2 Distributional effects of policy through the life-course**

This section provides further detail concerning the features underlying the projected aggregate effects of policy counterfactuals discussed in Section 3.1, by exploring the distributional variation of effects projected through the life-course. Discussion focusses on projections for families with individuals who were born between 1981 and 1990. Averaging over 10 birth cohorts dampens statistical noise associated with smaller samples, and the birth cohorts singled out here were aged between 26 and 35 in 2016, so that their projections capture the influence of the policy counterfactuals throughout the adult lifetime. Statistics calculated for the effects of the 10% rise in all income tax rates, distinguished by lifetime income quintile and cohort member age bands are reported in Table 3.

One of the clearest features that is evident in the statistics reported in Table 3 is the extent to which responses are skewed toward the upper end of the (lifetime income) distribution under all three behavioural scenarios. This is more than a proportional reflection of the scale of financial disparities described by the distribution, as is made



**Table 3:** Projected effects of a 10% rise in all income tax rates on average finances of families with members born between 1981 and 1990, distinguished by lifetime income quintile, age of cohort member, and accommodated behavioural response.

response	no behavioural responses			labour supply			labour supply and savings		
quintile	lowest	middle	highest	lowest	middle	highest	lowest	middle	highest
age-band	labour time per adult per week (mins)								
25-44	-14	-22	-16	-2	-7	-12	-2	-6	-3
45-54	-7	-38	-29	-1	-10	-15	-1	-4	3
55-64	-3	-21	-22	-1	-4	-11	-1	0	-3
65-74	-1	-13	-15	0	-3	-5	0	-3	1
75-84	0	0	0	0	0	0	0	0	0
consumption expenditure (£ per week)									
25-44	-2	-22	-54	-2	-20	-52	-1	-11	-32
45-54	-1	-23	-87	0	-19	-83	0	-16	-58
55-64	0	-14	-80	0	-9	-73	0	-11	-57
65-74	-1	-11	-72	-1	-5	-64	-1	-8	-57
75-84	-1	-9	-63	-2	-4	-53	-1	-6	-59
disposable income (£ per week)									
25-44	-5	-24	-72	-1	-19	-69	-1	-23	-63
45-54	-2	-34	-148	0	-21	-131	-1	-29	-119
55-64	-1	-28	-140	0	-15	-111	0	-17	-106
65-74	-5	-31	-139	-5	-18	-109	-5	-15	-105
75-84	-10	-23	-73	-10	-17	-72	-10	-10	-73
wealth (£ '000s)									
25-44	-1	-1	0	0	0	-1	0	-3	-6
45-54	-2	-13	-30	1	-8	-26	-1	-14	-35
55-64	-2	-26	-67	0	-15	-53	-1	-17	-59
65-74	-2	-32	-91	-1	-17	-74	-1	-21	-91
75-84	-7	-41	-118	-4	-22	-94	-5	-27	-121

Notes: See notes for table 1. Quintiles defined with respect to disposable family income, equivalised using the revised OECD scale, discounted, and aggregated over the life-course. All standard errors not greater than 1 unit of respective statistic. All financial statistics reported in 2016 prices. "disposable income" denotes private income net of government taxes and transfers. "wealth" denotes the aggregate of all private pension and non-pension wealth. Source: Author's calculations on simulated data generated using 30 separate sets of random draws.

clear by the effects on labour time reported at the top of the table. Rather, the result is a product of the fact that income taxes have a less pronounced bearing on families at the bottom of the distribution than they do at the top. Distributional variation of this sort is common, and is one of the key motivations originally put forward for the development of microsimulation modelling (e.g. Orcutt, 1957, pp. 116-117). It is also important for the current analysis, as it emphasises incentives of individuals toward the top of the distribution when considering the implications of the policy counterfactual on population aggregates, as discussed in Section 3.1.

Starting with the projected effects of the policy counterfactual on labour time reported at the top of Table 3, the differences in scale between alternative behavioural scenarios discussed in the preceding section are clearly evident. Results reported for the full-response scenario indicate that the reductions in labour supply described by the population aggregates late in the simulated time horizon are largely driven by families under the top income quintile. Of the three population subgroups reported in Table 3, the middle income quintile displays the most substantive reductions in labour supply under the full-response scenario, and then toward the start of the working lifetime. The relative scale of employment effects reported for the full-response scenario, across both age bands and income quintiles, reflect a balance between the price and wealth effects of the rise in income tax rates, where wealth effects are intensified with the projected decline in wealth holdings.

As discussed in Section 3.1, the wealth effects of the rise in tax rates are muted in the reduced-form descriptions for behaviour assumed here, relative to the fully-structural framework. This explains why the projected reductions in labour supply for the top quintile do not fall away under the non-response scenario, as discussed above for the full-response scenario. Furthermore, labour effects generated under the employment-response scenario lie between those projected under the non- and full-response scenarios. In this case, results of the employment-response scenario for the middle quintile are substantively closer to those of the full-response scenario, but for the highest quintile tend to be closer to the non-response scenario, reflecting distributional differences in scale of wealth effects. These differences consequently emphasise the crucial bearing of the specification used to identify behaviour using reduced-forms, which is an obvious concern wherever choice of specification is heuristically made with respect to limited objective criteria.

The projected declines in labour supply under the respective behavioural scenarios, combined with the rise in income tax rates, explain the projected declines in disposable income reported in the table. These, in turn, translate into the reported consumption effects. Notably, bearing in mind the differences in scale of projected labour supply effects, the effects on consumption and disposable income between the three behavioural scenarios appear broadly comparable. This is reflected in the close correspondence, across both age and income quintiles, projected for the effects on wealth by the three alternative behavioural scenarios. These results underscore the potential for well-specified reduced-form models to capture the qualitative nature of structural projections, even if differences in scale are almost impossible to avoid.

Table 4 reports statistics for the 20% fall considered for state retirement benefits that are similar to those described above for the rise in income tax rates. In common with the analysis of effects of the counterfactual increase in tax rates, the statistics reported in Table 4 reveal qualitatively similar projections under the three alternative behavioural scenarios, subject to noticeable differences in scale. All three behavioural scenarios suggest that households above the bottom quintile will tend to off-set the decline in state retirement benefits by increasing labour supply, reducing consumption, and increasing saving during the working lifetime. In contrast,

**Table 4:** Projected effects of a 20% fall in all state retirement benefits on average finances of families with members born between 1981 and 1990, distinguished by lifetime income quintile, age of cohort member, and accommodated behavioural response.

response	no behavioural responses			labour supply			labour supply and savings		
qunitile	lowest	middle	highest	lowest	middle	highest	lowest	middle	highest
age-band									
labour time per adult per week (mins)									
25-44	0	4	4	1	3	3	0	1	2
45-54	0	6	9	0	-1	1	0	2	3
55-64	0	11	35	0	1	7	0	5	11
65-74	0	-2	16	0	1	3	0	2	11
75-84	0	0	0	0	0	0	0	0	0
consumption expenditure (£ per week)									
25-44	-1	-7	-3	-1	-5	-2	-1	-1	-1
45-54	-1	-12	-5	-1	-12	-5	-1	-2	-1
55-64	-2	-8	-4	-2	-11	-11	-1	-4	-2
65-74	-24	-12	0	-24	-14	-14	-5	-11	-5
75-84	-46	-13	2	-46	-15	-15	-6	-16	-12
disposable income (£ per week)									
25-44	0	2	3	0	2	2	0	1	1
45-54	1	9	13	1	5	3	1	2	4
55-64	0	19	30	0	13	12	0	5	12
65-74	-52	-32	-7	-51	-34	-25	-52	-46	-24
75-84	-81	-71	-23	-81	-74	-57	-84	-84	-62
wealth (£ '000s)									
25-44	0	4	5	0	3	3	0	1	2
45-54	2	19	18	2	14	12	1	5	6
55-64	2	34	46	2	27	25	1	10	14
65-74	0	45	86	0	36	39	-4	9	24
75-84	-14	27	77	-14	17	26	-39	-21	4

Notes: See notes for table 3.

Source: Author's calculations on simulated data generated using 30 separate sets of random draws.

households at the bottom of the distribution, being liquidity constrained, do not increase provisions ahead of retirement, and are projected to suffer relatively large declines in consumption later in life as a result.

The adjustments discussed above are most substantive under the non-response scenario. This can be understood as the product of an excessive reflection of the wealth effect of the policy reform in the assumed reduced-form description for behaviour, relative to the fully-structural projection. The implication is that the reduced-form projections generate larger increases in wealth held late in life under both the non- and employment-response scenarios, relative to the full-response scenario. The rise in wealth tends to depress the increase in labour time projected late in life under the employment-response scenario, relative to the non-response scenario. The message running through these results is deceptively simple: A reduced-form structure is capable of capturing responses generated by a structural framework to the extent that the associated incentives are reflected by the assumed reduced-form behavioural specification. Unfortunately, identifying a well-specified reduced-form description for behaviour is a non-trivial task.

#### 4 SUMMARY AND CONCLUSIONS

This paper adds to the evidence base for choosing between alternative approaches for projecting decision making in a microsimulation context by exploring two key research questions: (i) How important is it to account explicitly for the behavioural trade-offs implied by policy counterfactuals in dynamic microsimulation projections; and (ii) to what extent will an explicit consideration of labour supply incentives alone reflect projections in which policy trade-offs concerning both the labour / leisure and consumption / savings margins are explicitly accounted for? The first of these questions provides a sense of the overall importance of structural behavioural responses when using a microsimulation model to analyse the effects of policy change, and the latter considers the practical importance of analytically convenient assumptions for structural decision making. The analysis focuses upon statistics projected using a common model structure, for the effects of two generic policy counterfactuals, with respect to three alternative sets of behavioural assumptions.

The three sets of behavioural assumptions reported here are designed to capture key features of common assumptions applied in the existing literature. The first is a scenario in which behaviour is projected on a functional description defined for the base policy context. This scenario is designed to approximate the traditional microsimulation approach of projecting behaviour based on reduced-form statistical descriptions estimated on historical survey data. The second scenario extends upon the first by building in structural responses to policy counterfactuals in relation to labour supply. This second scenario is designed to reflect the extension of structural methods within the contemporary microsimulation literature. Finally, the third behavioural scenario extends structural responses to counterfactual policy contexts, by considering both the labour/leisure and consumption/savings margins, as is common in agent based models, and a growing subset of the microsimulation literature. The structural framework considered for analysis projects behaviour as if it maximises expected lifetime utility, which has been the foundation of (dynamic) economic theory for the last half century.

Starting with a policy description designed to approximate taxes and benefits prevailing in the UK in 2016, the two counterfactuals consider a rise of 10 percentage points in all rates of income tax, and a reduction in the value of state retirement benefits of 20 percentage points. These policy counterfactuals are common in

the existing literature, either as a focus of stand-alone interest, or as potential policy adjustments to achieve defined budgetary objectives. Furthermore, the features of the two policy counterfactuals are complementary, presenting very different incentive effects, and different influences during the life course.

The analysis reveals that a model based on a well-specified reduced-form description of behaviour can produce a qualitative approximation of projections in which behaviour is based on a structural utility maximisation framework. Furthermore, extending a well-specified reduced-form model to accommodate structural responses for employment is likely to generate a close quantitative approximation of short-to-medium-term (up to 10 years in the current study) projections derived from a model in which both employment and savings are projected as though to maximise expected lifetime utility. Nevertheless, important quantitative differences are likely to emerge in longer-term projections between models distinguished by whether the consumption/savings margin is structurally simulated. This is because of the difficulty associated with projecting similar quantitative measures for consumption using reduced-form and structural approaches, with any differences accumulating in the projections for household wealth, which in turn distort other margins of decision making (including employment). This is an important observation, given the cursory consideration that is sometimes paid to reduced-form descriptions for consumption when accommodating structural labour-supply decisions.

The analysis emphasises that a ‘well-specified’ reduced-form specification capable of reflecting the implications of a structural framework requires more than a close reflection of behaviour observed in a given policy context. It must also be adapted to reflect the bearing of considered policy counterfactuals on incentives. This is a difficult task, particularly where forward-looking expectations are likely to be important. While informal methods might be used to identify pertinent incentives, selecting an empirical specification that will match the implications of a structural framework is non-trivial. In this sense, the results can be interpreted as providing support for the recent research interest in identifying heuristic decision rules, associated with Agent-based economics.

Furthermore, the analysis highlights the varying demands of different analytical approaches; a behavioural description that works well as part of a structural framework, can result in pronounced mis-specification if applied (without adjustment) as a reduced-form. This observation suggests that models are best understood as constellations of assumptions that work in concert with one another, calling into question the common practice of adopting model assumptions from third-party studies with little discussion about the similarities/differences in the methodological approaches employed.

## ACKNOWLEDGEMENTS

I should like to thank participants in the 6th World Congress of the International Microsimulation Association (Torino, 2017), two anonymous referees, and the journal editor for useful comments on earlier drafts. The usual disclaimer applies.

## REFERENCES

Ball, R. J. (1963). The Cambridge Model of Economic Growth. *Economica*, 30, 180-190.

- Blundell, R., & Dias, M. C. (2009). Alternative approaches to evaluation in empirical microeconomics. *Journal of Human Resources*, 44, 565-640.
- Browning, M., & Lusardi, A. (1996). Household saving: micro theories and macro facts. *Journal of Economic Literature*, 34, 1797-1855.
- Campbell, D. (1979). Assessing the impact of planned social change. *Evaluation and Program Planning*, 2, 67-90.
- Conant, R., & Ashby, W. R. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1, 89-97.
- Dawid, H., Harting, P., van der Hoog, S., & Neugart, M. (2016). *A heterogeneous agent macroeconomic model for policy evaluation: Improving transparency and reproducibility*. (Bielefeld University Working Paper 06-2016)
- Dhaene, G., & Barten, A. P. (1989). When it all began: The 1936 Tinbergen model revisited. *Economic Modelling*, 6, 203-219.
- Goodhart, C. (1984). Problems of Monetary Management: The UK Experience. In C. Goodhart (Ed.), *Monetary Theory and Practice: The UK experience*. London: Macmillan.
- Gourinchas, P., & Parker, J. (2002). Consumption over the Life Cycle. *Econometrica*, 70, 47-89.
- Gustman, A. L., & Steinmeier, T. L. (2005). The social security early entitlement age in a structural model of retirement and wealth. *Journal of Public Economics*, 89, 441-464.
- IFS. (2015). *Green Budget: 2015*. (Institute for Fiscal Studies)
- Johansen, L. (1963). A multi-sectoral study of economic growth: Some comments. *Economica*, 30, 174-176.
- Keynes, J. M. (1935). Official Papers: The League of Nations, Professor Tingenen's Method. *Economic Journal*, 49, 558-577.
- Klein, L. R. (2004). The contribution of Jan Tinbergen to economic science. *De Economist*, 152, 155-157.
- Kydland, F. E., & Prescott, E. C. (1982). Time to build and aggregate fluctuations. *Econometrica*, 1982, 1345-1370.
- Li, J., & O'Donoghue, C. (2013). A survey of dynamic microsimulation models: uses, model structure and methodology. *International Journal of Microsimulation*, 6, 3-55.
- Lucas, R. E., & Prescott, E. C. (1971). Investment under uncertainty. *Econometrica*, 39, 659-681.
- Lucas, R. J. (1976). Econometric policy evaluation: a critique. In K. Brunner & A. Meltzer (Eds.), *The Phillips Curve and Labor Markets*. Amsterdam: North-Holland.
- Nelson, C. R. (1972). The prediction performance of the FRB-MIT-PENN model of the U.S. economy. *American Economic Review*, 62, 902-917.
- Orcutt, G. (1957). A new type of socio-economic system. *Review of Economics and Statistics*, 58, 773-797.
- Pylkkanen, E. (2002). Modelling wages and hours of work. In E. Pylkkanen (Ed.), *Studies on household labor supply and home production*. Gothenburg: Goteborgs Universitet.
- Richiardi, M. (2014). The missing link: AB models and dynamic microsimulation. In S. Leitner & F. Wall (Eds.), *Artificial Economics and Self Organization: Agent-Based Approaches to Economics and Social Systems*. Cham: Springer International Publishing.
- Rust, J. (2008). Dynamic Programming. In S. N. Durlauf & L. E. Blume (Eds.), *New Palgrave Dictionary of Economics*. New York: Palgrave Macmillan.
- Solow, R. M. (2004). Progress in economics since Tinbergen. *De Economist*, 152, 159-160.
- Tran, C. (2016). Fiscal policy as a temptation control device: Savings subsidy and social security. *Economic*

*Modelling*, 55, 254-268.

van de Ven, J. (2017a). Parameterising a detailed dynamic programming model of savings and labour supply using cross-sectional data. *International Journal of Microsimulation*.

van de Ven, J. (2017b). SIDD: An adaptable framework for analysing the distributional implications of policy alternatives where savings and employment decisions matter. *Economic Modelling*, 63, 161-174.

## SUPPLEMENTARY MATERIAL

The Appendix is available online at the IJM website.

## NOTES

<sup>1</sup>The first sets of National Accounts were published immediately following the end of the Second World War, with the UK publishing in 1946 (covering the period 1938 to 1945), and the US in 1947.

<sup>2</sup>Tinbergen's first model was published in 1936 for the Netherlands, and used 24 equations to relate 31 variables. See, e.g. Dhaene and Barten (1989).

<sup>3</sup>The forecasting performance of econometric models was also brought into question with the finding that such models often failed to improve on forecasts derived using simple extrapolations of the historical time series (see, e.g. Nelson, 1972).

<sup>4</sup>Early examples of CGE models built for developed economies include the MSG model for Norway (1960; Johansen, 1963), and the Cambridge Growth Project for the UK (from the 1960s; Ball, 1963). Kydland and Prescott (1982) is most commonly cited as the first study to consider a DSGE model, although related models emerged in the early 1970s (e.g. Lucas and Prescott, 1971).

<sup>5</sup>One potential difficulty associated with ABMs is that, depending upon the care taken to set up the decision-making heuristics and learning rules, such models can remain exposed to the Lucas critique for policy evaluation purposes.

<sup>6</sup>Impact effects are implicitly defined here as effects that do not take account of the influence on incentives of a policy reform.

<sup>7</sup>Obtained by equating  $y = \alpha + A + (1 - t)(x_i^0 - B)$  to  $y' = \alpha + (1 - \Delta t)A + (1 - t - \Delta t)(x_i - B)$ , accounting only for income taxes, and ignoring investment income and potential transitions between marginal tax regimes.

<sup>8</sup>If the (imputed) annuity was insufficient to off-set the full decline in disposable income, then any shortfall was ignored for coordinates that had not yet accessed pension wealth, or deducted from projected consumption (on the assumption of liquidity constraint).

<sup>9</sup>Appendix B reports sensitivity of results to an alternative description for reduced-form behaviour, which projects no change in employment under the higher tax regime in 2016. This alternative projects a rise in net government revenue of £52 billion in 2016, which is roughly consistent with the HMRC projection; see HMRC publication "Direct effects of illustrative tax changes". Although past versions of the publication do not appear to be publicly available, the figure quoted here is also reported in IFS (2015), Chapter 10.

<sup>10</sup>This is because of the importance of labour income in determining taxable income, especially in the short-term where savings (and investment income) are very similar under alternative behavioural scenarios.