



Exploring the importance of behavioural endogeneity for policy projections

J. van de Ven

NIESR, 2 Dean Trench Street, London SW1P 3HE, UK and MIAESR, The University of Melbourne, Australia
justin.ven@simdynamics.org

ABSTRACT: Behavioural endogeneity is appealing functionality for any analytical tool designed to explore the implications of public policy alternatives. This study improves the evidence base for choosing between alternative approaches for projecting decision making by exploring two key research questions: (i) how important are the impact effects of policy change, relative to associated incentive effects; and (ii) to what extent can the over-all influence of behavioural responses to policy change be approximated by labour supply responses alone? Results obtained highlight the importance of selecting an analytical tool that is tailored to the subject of interest, and provide support for the view that best-practice evaluation of policy counterfactuals should consider the sensitivity of key results to alternative analytical approaches.

KEYWORDS: Dynamic microsimulation, Policy evaluation, Savings behaviour, Labour Supply

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, related interest has focussed on introduction into microsimulation models of structural methods for projecting agent decisions, motivated by the objective to obtain a conceptually coherent basis for exploring the effects of policy counterfactuals. Despite improvements in the analytical tools available to model builders, however, the 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 exploring 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 shift in focus in economics in favour of mathematical methods, in the form of 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.¹ The start of modern economic modelling can be traced to attempts to make use of the newly available data via the then newly developed economic 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.² 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 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 complex 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, after Lucas' critique of the validity 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, after unemployment and inflation were observed to increase in tandem.³ In response to this critique, the focus of economic model development shifted from projecting decision making 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 development of Computable General Equilibrium (CGE) models from the early 1960s, and Dynamic Stochastic General Equilibrium (DSGE) models from the 1980s.⁴

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 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 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 contexts. 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 population cross-section through time for the purposes of public 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).⁵

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 are the impact effects of policy change, relative to associated incentive effects? And secondly, to what extent can the over-all influence of behavioural responses to policy change be approximated by labour supply responses alone? The answer to the former of these questions provides a sense of the overall importance of behavioural responses when analysing policy change, and the latter indicates the practical importance of analytically convenient modelling assumptions that marginalise dynamic considerations 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) the functional description of behaviour remains unchanged, relative to the base policy environment; ii) labour supply responds optimally to the policy counterfactual, on the assumption that the functional description of savings remains as it was under the base policy environment; and iii) both labour supply and savings respond to the policy counterfactual to maximise expected lifetime utility.

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. This website also includes a set of video tutorials that walk-through the analysis, so that it should be possible to replicate all reported results.

An overview of LINDA is provided in Section 2.1, 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 described 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; extended discussion of the model structure is provided in (van de Ven, 2017b) and details concerning the model's parameterisation to UK data are reported 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 ^u |
| - number of dependent children ^u | - age of dependent children ^u | - student status ^u |
| - education status ^u | - private pension wealth ^d | - timing of pension access ^d |
| - non-pension wealth ^d | - wage potential ^u | - immigration ^u |
| - emigration ^u | - survival ^u | |

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. This 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.

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; λ^o is a random adjustment factor that is included to allow for involuntary unemployment (lack of a wage offer); λ^{emp} adjusts to reflect the effect of labour supply decisions on earnings (varying with endogenous l); and λ^{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 λ^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 σ_ω^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 τ 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 equivalised (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. 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, $rr_{i,a}^h$; 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, π_{er} , or the employee, π_{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 taken to evaluate the effects of each policy counterfactual. The methods used to evaluate measures of uncertainty for the 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. The second 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 under all policy considered environments. State retirement benefits are defined here as the basic State Pension (a contributory flat-rate benefit), the Pension Credit (a means-tested retirement benefit), and the personal allowances for Housing and Council Tax benefits (used to evaluate means-testing of each benefit). 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 2046 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 the

random draws that were used to project individual characteristics that are considered to evolve with uncertainty from 2016 to 2046 under the base policy scenario. All simulated data – other than the random draws – projected for the period 2017 to 2046 under the base scenario were then 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.

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 disparities affecting its projections.

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.

LINDA is designed so that each simulated adult is associated with the same weight in the respective population cross-section. The model parameterisation considered here assumes 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

The current analysis explores the influence of three alternative behavioural scenarios on the projected effects of policy counterfactuals: Behaviour that does not respond to the altered incentives associated with a policy counterfactual; labour supply that responds to altered incentives of a counterfactual but savings behaviour that does not; and the case in which both savings and labour supply respond to changes in policy incentives. These behavioural alternatives are interesting because they reflect common assumptions made in the contemporary microsimulation literature.

One of the most simple methods for projecting decisions is on the basis of reduced-form descriptions of past behaviour. In its most general form, the approach involves the assumption of a functional relation between the behaviour of interest and alternative ‘explanatory’ characteristics, and is consequently capable of capturing almost any conceivable relationship. The central problem with this most general form is the difficulty associated with identifying ‘appropriate’ assumptions to adopt for the bearing of explanatory variables on a given behavioural margin. The most common approach for resolving this problem is to estimate an hypothesised relationship econometrically, so that it describes correlations in historical survey data. This approach requires that only relationships between observable characteristics are included for analysis. Despite this limitation, however, the approach retains a relatively high degree of flexibility and simplicity, which is why it has featured prominently in microsimulation models since their first inception.

The reduced-form approach for projecting behaviour is reflected in this analysis by the assumption that behaviour ‘does not respond to the altered incentives associated with a policy counterfactual’. Note that this assumption does not imply that projected behaviour will be invariant in context of a policy counterfactual, as projected responses will be generated wherever a counterfactual influences the ‘explanatory characteristics’ assumed for describing behaviour. In general, these characteristics may explicitly include policy parameters of interest; a reduced-form description of employment participation, for example, might include (observable) marginal tax rates in its set of assumed explanatory variables. Any change in these tax rates would then imply a change in projected employment participation.

In contrast to the general reduced-form description noted above, however, the current analysis limits the set of explanatory variables to the 14 individual-specific characteristics listed at the top of Section 2.1. One implication of this restriction is that fiscal policy can influence projected decisions under the ‘non-response’ behavioural scenario only indirectly, through its bearing on accumulated wealth balances. This limitation is imposed to ensure that all of the behavioural alternatives considered for analysis share a common analytical approach, thereby facilitating comparisons between them.

The conceptual flexibility of the reduced-form method for projecting behaviour is both its principal strength, and source of weakness. On the one hand, the approach’s flexibility makes it possible to reflect a very wide range of potential relationships in a behavioural projection. On the other, its indeterminacy raises difficult questions concerning the most appropriate relationships to include for analysis. One of the key objectives when choosing a reduced-form specification for projecting behaviour is that relationships implied by the specification

should be stable in context of a changing policy environment. It is important to recognise that the ability to achieve this objective is constrained here by the limitation that is imposed on the set of explanatory variables for behaviour. Nevertheless, even in an unconstrained context, the difficulties associated with defining a reliable reduced-form specification are well recognised, which is an important motivation for the development of structural approaches for projecting behaviour.

It is generally useful when implementing structural descriptions for behaviour in a microsimulation context to adopt assumptions that marginalise consideration of intertemporal dynamics. This is because allowing for temporal dynamics generally results in a substantive increase in the complexity of the utility maximisation problem. A common simplifying assumption, for example, is to exogenously set discretionary consumption equal to disposable income when exploring structural models of labour supply; it is this type of assumption that the intermediate behavioural alternative considered here is designed to reflect. The last of the three behavioural alternatives considered for analysis complete the spectrum by using structural methods to project joint savings and labour supply decisions.

2.3.1. Behavioural projections 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 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 12 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 accom-

modated 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 size 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 current analysis.

2.3.2. Adapting LINDA to explore behavioural alternatives

LINDA is designed to consider the case in which both savings and labour supply respond to changes in policy incentives, thereby capturing one of the three behavioural alternatives by default. This is achieved by taking account of the simulated policy environment in the forward projections of individual specific characteristics when expected lifetime utility is evaluated – denoted (ii) in the above discussion.

The assumption that the functional description of behaviour remains unchanged, relative to the base policy environment, required no adjustment of LINDA's structure to explore. All that was required was to load in the

grid solution obtained under the base policy environment, and use this for projecting the population forward through time under the policy counterfactual – essentially skipping over the first stage of a standard simulation procedure.

In contrast, 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 adjustment of the existing model code. In this case, the model was adapted to first load in the results of the grid evaluated under the base policy scenario. The numerical routines that search for an optimum to the utility maximisation problem were then altered to hold consumption and pension savings fixed at their loaded values, and to search only over the discrete labour supply alternatives considered for analysis (see Section 2.1). The results of this constrained optimisation were saved by the model, after which each simulation proceeded as normal.

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. Projections for the rise in income tax rates are discussed first, before addressing the projected effects for the fall in benefit values.

The rise in tax rates is projected, under the assumption of no behavioural responses to the counterfactual (the non-response scenario), to increase net government revenue by £52 billion in 2016, rising to £81 billion by 2046. The projected budgetary effect for 2016 is similar to statistics 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.⁶ Importantly, the projected impact of the rise in income tax rates reported by the HMRC exogenously assumes no behavioural responses to adjustments under the higher rate of tax, and that for higher rate taxpayers a 1 per cent rise in the marginal tax rate will lead to a 0.45 per cent reduction in taxable incomes. We return to discuss this point in the following subsection.

Comparing the net budgetary effects reported in Table 1 for the rise in income tax rates projected under the non-behavioural scenario against those projected under the assumption that both savings and labour supply respond to the counterfactual (the full-response scenario) reveals a high degree of sensitivity to alternative behavioural assumptions. The table indicates that extending the non-response scenario to accommodate labour and savings responses to the rise in income tax rates depresses the projected increase in net government revenue by approximately 20% (£10 billion) in 2016, widening to 30% (£26 billion) by 2046 (relative to the assumption of

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

response*	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	52.4	45.8	41.8	17.5	17.9	17.2	0.1	0.6	0.6	-19.6	-19.6	-19.7
	(0.4)	(0.4)	(0.4)	(0.1)	(0.1)	(0.1)	(0.0)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
2021	56.5	50.7	44.4	19.3	19.8	19.3	0.4	0.8	0.6	-22.0	-22.1	-22.3
	(0.5)	(0.5)	(0.4)	(0.1)	(0.1)	(0.1)	(0.0)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
2026	59.4	53.1	45.2	21.8	22.4	22.2	0.6	1.0	0.7	-25.1	-25.2	-25.5
	(0.7)	(0.7)	(0.6)	(0.1)	(0.1)	(0.1)	(0.0)	(0.1)	(0.1)	(0.1)	(0.1)	(0.1)
2036	58.1	50.2	39.3	30.8	31.7	32.1	9.3	9.7	9.0	-35.4	-35.5	-36.2
	(0.8)	(0.8)	(0.6)	(0.2)	(0.2)	(0.2)	(0.3)	(0.3)	(0.3)	(0.2)	(0.2)	(0.2)
2046	81.2	70.1	55.4	42.0	43.4	44.3	1.6	1.9	0.8	-47.6	-47.7	-48.8
	(0.9)	(0.8)	(0.7)	(0.2)	(0.2)	(0.3)	(0.1)	(0.1)	(0.1)	(0.2)	(0.2)	(0.2)
	income tax revenue						consumption tax revenue					
2016	52.3	46.3	45.1	-2.1	-1.7	-1.4	0.1	0.1	-2.7	0.0	0.0	-1.1
	(0.4)	(0.4)	(0.4)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
2021	58.0	52.9	48.7	-2.5	-2.2	-1.9	-1.1	-1.3	-3.6	-0.1	-0.1	-1.1
	(0.5)	(0.5)	(0.5)	(0.0)	(0.0)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
2026	61.8	56.3	50.1	-3.0	-2.5	-2.2	-1.8	-2.2	-4.2	-0.3	-0.3	-1.1
	(0.7)	(0.7)	(0.6)	(0.0)	(0.0)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
2036	70.1	63.2	53.7	-4.0	-3.3	-2.8	-2.7	-3.3	-5.4	-0.6	-0.5	-1.3
	(0.7)	(0.7)	(0.5)	(0.0)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)
2046	85.7	75.7	62.5	-4.8	-3.6	-3.1	-2.9	-3.8	-6.4	-0.8	-0.7	-1.4
	(1.0)	(0.8)	(0.7)	(0.0)	(0.1)	(0.1)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)	(0.0)

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

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.

no responses). This sensitivity of projected budgetary effects to alternative behavioural assumptions emphasises the potential importance of the approach taken to project decisions, particularly in a dynamic context.

The remaining three panels of Table 1 disaggregate the projections for net government revenue, providing detail that helps to explain the variation between behavioural alternatives that is noted above. 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 non-response scenario, relative to the full-response scenario. That is, the higher increases in net government revenues projected under the assumption of no behavioural responses 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.⁷ 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) provide some support for this conjecture. Specifically, the projected impact of the rise in tax rates on net government revenues under the employment-response scenario are between those generated under the non- and full-response scenarios. In the near term, these intermediate projections are closer to those in which savings and labour supply respond optimally to the policy counterfactual, but drift toward those in which savings and labour decisions are insensitive to counterfactual incentives as the time horizon is extended. We return to discuss the reasons for this shift below.

Although the considered rise in income tax rates has a larger projected impact on net government revenues than the 20% fall in the value of retirement benefits reported in Table 1, both counterfactuals are shown to have substantive capacity to alter government finances. Furthermore, in contrast to the results reported for the rise in income tax rates, the projected effects on net government revenues of the fall in state retirement benefits are found to be broadly insensitive to alternative assumptions concerning behavioural responses. This is an important result, highlighting the potential usefulness of tailoring the features of considered policy reforms to the limitations of a microsimulation model (and vice versa). Put another way, the results suggest that altering the terms of retirement benefits is less behaviourally distortional – at least in terms of projections for net government revenue – than altering income tax rates.

While the effects of alternative behavioural assumptions on projected net government revenues are muted under the retirement benefits counterfactual, some interesting differences between the projections are discernable. The increase in net government revenues projected for 2016 is £300 million smaller 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 £2.3 billion higher under the full-response scenario. The remaining statistics in Table 1 indicate that £1.1 billion of this temporal reversal is due to a relative decline in state expenditure on retirement benefits under the full-response scenario, £1.0 billion is due to a smaller fall

of income tax revenue, and £500 million is due to a smaller decline in consumption tax revenue. The relative importance of expenditure on state benefits in explaining these fluctuations is consistent with the focus of the considered policy counterfactual. Further detail can be obtained by exploring the contemporaneous variation generated for household finances, to which we now turn.

Table 2 reports population aggregate effects on domestic sector financial aggregates. Starting with statistics reported for the effects of the rise in income tax rates, the results highlight the substantive differences in projected decision making associated with alternative behavioural assumptions. Under the non-response scenario, the rise in income tax rates is projected to have no immediate effect on labour supply, but does result in increasing employment as the time horizon is extended. By 2046, almost one million additional (equivalent) full-time workers are projected under the higher tax rate scenario when behavioural responses to altered incentives are suppressed (from a total population of 69 million adults). This increase in simulated labour supply with the projected time horizon is a product of the coincident decline in (non-pension) wealth. The rise in income tax rates from 2016 is projected under the non-response scenario to result in a decline in aggregate net wealth by 2046 of £775 billion (885 - 110).

The increase in private income associated with the projected rise in labour supply is more than sufficient to off-set the coincident decline in investment income associated with falling net assets, so that private income is projected to increase by £30 billion in 2046 under the high tax rate counterfactual when behavioural responses are suppressed. This increase in private income amplifies the increase in income tax revenue associated with the higher tax rates that are imposed. It also mitigates the projected impact on disposable incomes, thereby supporting a higher consumption stream. Furthermore, the increase in employment rates permits the accrual of higher pension wealth, off-setting the overall decline in total net worth associated with the higher tax rate counterfactual.

Interpreted from a utility maximising perspective, the projections that omit behavioural responses to altered incentives assume that the population spends and supplies labour under the counterfactual policy environment as though they still lived under the base policy environment. This means that, in context of the 10% rise in income tax rates, people are projected to spend as though they enjoyed a higher disposable income than they actually receive for any given measure of private income. Conceptually, these individuals are repeatedly “surprised” by how little wealth they have in each succeeding period, and subsequently cut back consumption and increase labour supply to make up the ever-growing short-fall in wealth. The result is a dynamic projection for the counterfactual that reflects a run-away wealth effect.

Results reported for the full-response scenario to the rise in income tax rates contrast starkly with those in which behavioural responses are suppressed (discussed above). The top-left panel of Table 2 indicates that, whereas the non-response scenario generates no effect following the rise in income tax rates on labour supply in the short-term and a substantive increase in the longer term, the full-response scenario projects a substantive decline in labour in the short term that declines with the simulated time horizon. In 2016, the number of equivalent full-time workers is projected under the full-response scenario to fall following the rise in income tax rates by 570,000, reducing to 150,000 by 2046, by which time the difference with the labour effects projected under the non-response scenario is over one million effective full-time workers.

The falls in labour supply generated under the full-response scenario reflect the reduction in the effective price

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

response*	income tax rates			retirement benefits			income tax rates			retirement benefits		
	rise 10%			fall 20%			rise 10%			fall 20%		
	none	emp	full	none	emp	full	none	emp	full	none	emp	full
	equivalent full time employees ('000)						consumption expenditure					
2016	-0.1	-515.3	-573.4	0.0	90.6	154.8	0.0	-0.1	-18.7	0.0	-0.2	-7.4
	(0.0)	(9.1)	(9.6)	(0.0)	(5.1)	(6.4)	(0.0)	(0.0)	(0.2)	(0.0)	(0.0)	(0.0)
2021	298.1	-137.0	-343.0	10.3	96.4	141.2	-8.2	-9.7	-24.8	-1.3	-1.3	-7.6
	(6.4)	(14.1)	(13.4)	(1.2)	(7.0)	(7.3)	(0.1)	(0.1)	(0.2)	(0.1)	(0.1)	(0.1)
2026	506.5	82.5	-214.9	14.4	107.5	148.1	-13.0	-15.4	-29.0	-2.5	-2.3	-8.0
	(10.6)	(13.1)	(12.3)	(1.8)	(6.9)	(7.1)	(0.2)	(0.2)	(0.2)	(0.1)	(0.1)	(0.1)
2036	825.9	351.6	-106.3	11.3	133.2	171.6	-18.9	-23.2	-36.9	-4.8	-4.4	-9.2
	(13.9)	(16.3)	(15.0)	(1.5)	(6.7)	(7.9)	(0.2)	(0.2)	(0.3)	(0.1)	(0.1)	(0.1)
2046	940.6	433.1	-147.6	10.7	167.0	185.2	-20.7	-27.2	-44.4	-7.1	-6.1	-10.4
	(14.1)	(17.2)	(15.8)	(0.8)	(9.4)	(12.3)	(0.3)	(0.3)	(0.3)	(0.2)	(0.2)	(0.2)
	private income						disposable income					
2016	0.0	-15.1	-17.1	0.0	1.5	3.2	-52.3	-59.8	-61.8	-17.5	-16.6	-16.0
	(0.0)	(0.4)	(0.4)	(0.0)	(0.1)	(0.1)	(0.4)	(0.5)	(0.5)	(0.1)	(0.1)	(0.2)
2021	7.2	-6.3	-12.8	-1.1	0.3	2.8	-50.9	-57.9	-61.2	-20.7	-19.8	-18.5
	(0.4)	(0.6)	(0.6)	(0.0)	(0.1)	(0.2)	(0.6)	(0.6)	(0.6)	(0.2)	(0.2)	(0.3)
2026	12.3	-3.0	-12.3	-2.2	-0.2	3.3	-50.0	-58.2	-62.5	-24.3	-23.1	-20.9
	(0.6)	(0.7)	(0.6)	(0.0)	(0.2)	(0.3)	(0.5)	(0.7)	(0.9)	(0.3)	(0.3)	(0.3)
2036	22.7	1.7	-14.2	-4.5	-1.0	4.8	-40.4	-52.5	-60.2	-36.0	-33.5	-29.8
	(0.9)	(1.1)	(0.8)	(0.1)	(0.2)	(0.3)	(0.7)	(0.8)	(0.8)	(0.4)	(0.5)	(0.4)
2046	30.4	-0.2	-23.4	-7.1	-1.1	6.4	-56.2	-74.5	-85.9	-50.0	-45.7	-40.6
	(1.2)	(1.1)	(1.1)	(0.1)	(0.3)	(0.5)	(1.6)	(1.5)	(1.7)	(0.4)	(0.5)	(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	1.7	-21.6	-7.5	0.1	3.8	17.4	-233.2	-259.3	-175.6	-83.7	-81.0	-39.3
	(0.7)	(1.8)	(1.9)	(0.1)	(0.8)	(1.1)	(2.4)	(3.1)	(2.5)	(0.7)	(1.0)	(1.5)
2026	11.1	-33.1	-5.4	0.2	7.0	32.4	-420.3	-473.5	-326.7	-171.7	-165.0	-79.9
	(1.3)	(2.1)	(2.2)	(0.2)	(1.1)	(2.2)	(4.2)	(4.3)	(3.5)	(1.7)	(1.7)	(2.2)
2036	53.9	-39.0	27.8	0.4	13.3	72.4	-681.3	-799.3	-566.0	-364.4	-344.6	-173.4
	(3.3)	(4.2)	(3.6)	(0.2)	(1.2)	(2.4)	(5.3)	(6.8)	(6.2)	(3.4)	(3.8)	(4.2)
2046	110.5	-44.2	81.4	0.6	23.8	136.2	-885.0	-1108.2	-793.1	-595.2	-549.5	-286.3
	(3.9)	(4.6)	(4.0)	(0.3)	(1.7)	(4.0)	(7.8)	(9.4)	(9.7)	(4.4)	(5.3)	(4.8)

Notes: See notes for table 1. "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.

of leisure implied by higher income tax rates. The influence of this price effect on simulated projections is dampened as the time horizon is extended, due to the exaggeration of wealth effects associated with the coincident decline in the value of net assets held. Optimising agents are projected to substitute out of discretionary consumption and into leisure, sacrificing disposable income and drawing down their wealth balances.⁸ Seen from this perspective, differences between the projected effects of the rise in income tax rates generated under the non-behavioural and full-behavioural scenarios are primarily attributable to the fact that the price and income effects of the counterfactual work in opposite directions: Whereas the non-response scenario reflects a runaway wealth effect, it is the price effect that dominates projections under the full-response scenario.

The labour supply response to the rise in income tax rates projected for 2016 under the employment-response scenario, closely reflects that generated under the full-response scenario. Despite the coincident decline projected for disposable incomes in 2016 (due to both the reduction in labour incomes and higher tax burdens), however, projected consumption expenditure remains invariant by construction under the employment-response scenario. This combination of lower disposable incomes and fixed consumption expenditure in near-term projections following the rise in income taxes produces a rapid draw-down of assets under the employment-response scenario. The decline in asset holdings exaggerate the wealth effects of the higher tax rate counterfactual, offsetting incentives to reduce labour supply under the employment-response scenario. This is the mechanism that is responsible for the drift of the projections under the employment-response scenario toward those of the non-response scenario as the time horizon is extended.

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. Furthermore, whereas the rise in income tax rates directly affect individuals throughout their adult lives, the fall in state retirement benefits only has a direct impact on individuals beyond state pension age. In context of a preference relation that includes temporal discounting in evaluating forward looking expectations, both of these considerations tend to dampen differences in projected employment effects of the reduction in state retirement benefits between the non-response and a full-response behavioural scenarios. These factors are consequently responsible for the broad similarities in the effects on the evolving government budget projected for the fall in retirement benefits based on all three behavioural scenarios.

Relative to the non-behavioural response scenario, the full-response scenario projects larger increases in labour supply throughout the simulated horizon. This dampens the decline in income taxes associated with the coincident fall in the value of (taxable) state retirement benefits. In contrast, the full-response scenario projects larger declines in consumption, due to the fall in lifetime (disposable) incomes following the counterfactual, which reduces aggregate consumption taxes. This disparity is mitigated with the simulated time horizon due to the relative decline in wealth under the non-behavioural scenario. The relative decline of wealth projected following the fall in state retirement benefits is of particular note: In contrast to the behavioural insensitivity of projections for net government revenues, projected effects of the fall in state retirement benefits on private wealth balances diverge substantively between the behavioural scenarios with the simulated time horizon. This observation underscores the fact that, whereas alternative behavioural assumptions may have a muted bearing on one margin of interest (e.g. the net government budget), they can project substantive differences on other margins (e.g. measures of household sector wealth).

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. This is more than a proportional reflection of the scale of financial disparities described by the distribution, as is made clear by the effects on labour time reported in 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 top-left panel of Table 3, employment effects under the non-response scenario are projected to increase substantively with lifetime income quintile. In this case, the correlation between lifetime income and employment effect is due to the distributional bearing that the increase in income tax rates has on wealth holdings, as reported in the bottom-left panel of the table. The reason for this fall in wealth under the non-response scenario is made clear by the middle two panels of the table: Despite the increases in labour supply projected for the middle and highest income quintiles, disposable income continues to fall under the high tax rates counterfactual, and the associated decline in consumption is insufficient to compensate.

As noted in Section 3.1, the relationships described above are the product of projecting behaviour on a functional description of decision making conceptually observable under the base policy environment. The projected effects can consequently be understood as the product of decisions that fail to account for the counterfactual increase in tax rates. While it might be possible to mitigate the influence of omissions like this on projections based on reduced-form descriptions of decision making through careful specification selection, it is difficult (if not impossible) to predict the out-of-sample success in this regard. It is precisely this problem that has motivated the development of structural approaches for projecting decisions in counterfactual contexts.

Turning to statistics reported for the full-response scenario in the right most panels of Table 3, some interesting distributional variation is evident. The effects on labour supply reported in the top-right panel of the table indicate that the reductions in labour supply that are described by the population aggregates discussed in Section 3.1 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-responses scenario, and then toward the start of the working lifetime. The relative scale of employment effects

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	0	3	16	-2	-4	8	-2	-6	-3
45-54	0	10	45	-1	1	26	-1	-4	3
55-64	0	10	45	-1	3	18	-1	-1	-2
65-74	0	3	39	0	-2	17	0	-3	2
75-84	0	0	1	0	0	0	0	0	0
consumption expenditure (£ per week)									
25-44	0	-5	-9	0	-6	-10	-1	-12	-32
45-54	0	-14	-25	0	-15	-30	0	-16	-58
55-64	0	-11	-26	0	-13	-37	0	-11	-57
65-74	0	-8	-29	0	-10	-48	0	-7	-56
75-84	0	-7	-40	0	-9	-64	0	-6	-58
disposable income (£ per week)									
25-44	-2	-21	-47	-2	-23	-54	-2	-23	-64
45-54	-1	-27	-74	-1	-30	-97	-1	-28	-119
55-64	0	-14	-54	0	-20	-90	0	-17	-106
65-74	0	-13	-51	0	-19	-95	0	-15	-103
75-84	-1	-9	-49	-1	-13	-73	-1	-10	-69
wealth (£ '000s)									
25-44	-1	-6	-11	-1	-6	-13	0	-3	-6
45-54	-1	-19	-38	-2	-23	-53	-1	-13	-35
55-64	-1	-20	-41	-2	-26	-76	-1	-17	-58
65-74	-1	-22	-56	-1	-30	-112	-1	-20	-89
75-84	-2	-28	-81	-2	-38	-144	-1	-27	-117

Notes: See notes for table 1. Quintiles defined with respect to disposable family income, equalised 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.

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.

The projected effects on labour supply discussed here are in sharp contrast to the assumptions that HMRC uses to evaluate the impact of changes in marginal tax rates. As discussed in Section 3.1, HMRC projections assume no responses to changes in rates under the highest rate of income tax, and that a 1% increase in the highest rate of tax will reduce taxable income by 0.45 percentage points. While the population aggregate statistics reported in Section 3.1 do support the conjecture that price effects will dominate income effects associated with the rise in income tax rates (consistent with the HMRC assumptions), the statistics reported here suggest that this is driven in the forward projections by responses of middle income families. This apparent disparity is attributable to the time-horizon considered for analysis. As noted in Section 3.1 wealth effects are exaggerated as the time horizon is extended, as consecutive periods under the high-tax environment result in lower accumulation of family wealth, particularly among high income families. This tends to exaggerate the bearing of price effects in the short- relative to the long-run. Results of this sort consequently serve to highlight the fact that the short-run effects of a policy counterfactual can provide a poor guide to longer-term implications.

The effects of the rise in income tax rates projected under the employment-response scenario that are reported in the middle panel of Table 3 generally fall between the non-response and full-response effects that are discussed above. The exception in this regard are the wealth statistics that are reported at the bottom of the table, and which indicate more substantive falls in asset holdings under the employment-response scenario than either of the behavioural alternatives. The relatively large declines projected for wealth under the employment-response scenario also result in more substantive declines in consumption and disposable income late in the life-course for the middle and highest lifetime income quintiles than under either of the alternative behavioural scenarios.

As discussed in Section 3.1, the effects described in the preceding paragraph are the result of the constrained optimisation that is imposed under the employment-response scenario, relative to either of the behavioural alternatives. The statistics reported for the full-response scenario indicate that, if given the choice, optimising families will generally trade consumption for leisure following the rise in income tax rates. Under the employment-response scenario, however, the consumption decisions of families are broadly constrained to reflect the non-response scenario. In this case the exchange of consumption for leisure is effectively achieved by adopting employment responses that are intermediate between the non-response and full-responses scenarios, accumulating fewer assets than in either of the behavioural alternatives, and consuming less in retirement as a result.

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. Starting with the left-hand side panels of the table, the reported statistics indicate that the fall in state retirement benefits is projected under the non-response scenario to have practically no impact during the working lifetime (under age 65), regardless of a family's financial circumstances. This feature of the results reflects similarly muted responses reported for employment aggregates in Section 3.1, and is a product of the fact that – unlike the counterfactual rise in tax rates – the reform to state retirement benefits does not directly affect individuals until they reach state pension age.⁹

Beyond state pension age, the results indicate that the fall in state retirement benefits is projected under all three behavioural scenarios to have a more substantive impact on families toward the bottom of the income distribu-

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		
quintile	lowest	middle	highest	lowest	middle	highest	lowest	middle	highest
labour time per adult per week (mins)									
25-44	0	0	0	0	1	0	0	1	1
45-54	0	0	0	0	2	1	0	2	3
55-64	0	0	0	0	6	9	0	5	11
65-74	0	0	4	0	3	16	0	2	12
75-84	0	0	0	0	0	0	0	0	0
consumption expenditure (£ per week)									
25-44	0	0	0	0	0	0	-1	-1	-1
45-54	0	0	0	0	0	0	-1	-3	-1
55-64	0	0	0	0	0	1	-2	-4	-2
65-74	-7	-5	-4	-10	-4	1	-16	-11	-5
75-84	-5	-16	-18	-5	-15	-11	-12	-16	-12
disposable income (£ per week)									
25-44	0	0	0	0	0	0	0	1	1
45-54	0	0	0	0	1	2	0	2	4
55-64	-1	-1	-1	-1	1	9	-2	5	12
65-74	-52	-55	-45	-52	-52	-27	-53	-46	-23
75-84	-71	-91	-80	-71	-90	-74	-74	-83	-61
wealth (£ '000s)									
25-44	0	0	0	0	0	0	0	1	2
45-54	0	0	0	0	0	1	1	5	6
55-64	0	0	0	0	1	4	2	10	14
65-74	-6	-9	-8	-6	-6	4	-3	9	25
75-84	-36	-49	-40	-35	-45	-24	-29	-21	5

Notes: See notes for table 3.

tion between ages 65 to 74, and a larger effect on high income families thereafter. This reflects the capacity of high income families to off-set the effects of the policy counterfactual by drawing upon their private wealth, up until the time that their wealth is exhausted.

The middle and right-most panels of Table 4 indicate that allowing for endogenous behavioural responses to the fall in state retirement benefits generates projections in which private decisions off-set the policy counterfactual. These off-sets are reflected by higher measures of projected wealth and smaller declines in disposable income. In this regard, the scale of the off-sets tend to increase with lifetime income, and the effects generated under the employment-response scenario are intermediate to those generated under the non- and full-response scenarios.

Labour supply is projected to increase following the fall in state retirement benefits under both the employment- and full-response behavioural scenarios. Interestingly, the scale of labour supply effects of the fall in state retirement benefits projected under the employment- and full-response scenarios are of similar magnitude, both through the life course and across the income distribution. The most substantive labour responses under both behavioural scenarios are projected for the highest income quintile late in the working lifetime. Such effects are completely absent under the non-response scenario.

Consumption under state pension age is only projected to be affected by the fall in state retirement benefits under the full-response scenario. In this case, consumption is projected to fall under the counterfactual, thereby exaggerating the rise in private wealth made possible by the increase in labour supply. While the consumption effects that are reported for the full-response scenario in Table 4 are not very large, they are sufficient to open up a sizeable gap in the wealth effects projected under the employment- and full-response scenarios.

In summary, while allowing for employment decisions in isolation does help to explain some of the difference between the non- and full-response scenarios, it generally does not approximate the full difference. Furthermore, the results reported here highlight limitations that stem from the difficulties associated with capturing forward-looking responses when projecting the effects of policy counterfactuals based on reduced-form descriptions for behaviour. In the current context, these limitations are reflected by the substantive differences in the effects of the fall in retirement benefits on wealth projected under alternative behavioural assumptions: Whereas the full-response scenario suggests that families will off-set the effects of the policy counterfactual through increased labour supply and private savings during the working lifetime, the non-response scenario projects no effects prior to state pension age. In the current context the budgetary influences of these differences are muted, in part because labour supply and consumption responses projected the full-response scenario off-set one another, and in part because of their limited scale. It would, however, be difficult to know that this would be the case, without actually running the associated analysis.

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 are the impact effects of policy change, relative to associated incentive effects; and (ii) to what extent can the over-all influence of behavioural responses to policy change be approximated by labour supply responses alone? The first of these questions provides a sense of the overall importance of 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 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 of behaviour observed in the past. 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.

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.

A key conclusion of the analysis – responding to question (i) defined above – is that the bearing of alternative behavioural assumptions on the effects of a policy counterfactual vary with both the policy counterfactual and projected margin of interest. In the current context, for example, the projected increase in net government revenue in 2016 following the rise in income tax rates is projected to be 20% (£10 billion) lower under the scenario in which both savings and labour supply respond to the coincident change in incentives than when they do not. This difference increases to one third (£26 billion) when the projected time horizon is extended to 2046. In contrast, there is very little difference between the effects on net government revenue of the fall in state retirement benefits projected under the alternative behavioural scenarios throughout the projected time horizon considered here. Furthermore, this relationship is broadly reversed if the margin of interest is the impact on aggregate household sector wealth rather than the net government budget. These results highlight the importance of selecting an analytical tool that is tailored to the subject of interest (and vice versa).

Comparisons between the three alternative behavioural scenarios – responding to question (ii) above – reveal that projections that account only for labour supply responses to the incentives implied by policy counterfactuals generally fall between projections that omit any responses, and those that account for both savings and labour supply responses. In the current context, the labour-response projections are generally found to provide a closer reflection of the projections that account for both savings and labour responses, when the projected time horizon is short. For effects of the rise in income tax rates on net government revenues in 2016, for example, projections that account for labour supply responses to the incentives of the policy counterfactual are more than two thirds closer to those that accommodate both savings and labour responses than to those that

accommodate no such responses. However, this bias is found to decline with the projected time horizon, and by 2046 it is marginally reversed in favour of the non-response scenario.

In the current analysis, the drift of projections that accommodate labour supply responses to incentives away from those that accommodate both savings and labour supply incentives is attributed to persistent mismatch between the two behavioural scenarios in projected savings rates. This mismatch accumulates with the projected time horizon, in differences projected for accumulated wealth, which have an indirect effect on labour supply decisions. This form of disparity could potentially be mitigated via selection of a reduced-form description for savings that provides a better approximation of the underlying structural framework than is considered here. Unfortunately, the difficulties involved in achieving such an objective are well understood, which is a prime motivation for the development of structural approaches for projecting decision making. These observations support the proposition that, given contemporary analytical limitations, best practice in counterfactual policy evaluation involves comparing projections from alternative analytical approaches, to ensure that stylisations associated with a preferred methodology do not over-turn centrally important results.

REFERENCES

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

¹The 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

²Tinbergen'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).

³The 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).

⁴Early 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).

⁵One 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.

⁶Estimate reported in HM Revenue and Customs 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.

⁷This 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.

⁸Nevertheless, they do respond to the increased value of the tax shielding associated with pension saving in the high tax regime, as indicated by the rise in pension wealth with the projected time horizon despite the decline in both labour supply and non-pension wealth.

⁹State pension age is defined in Section 2.2.