Economic Modelling 63 (2017) 161–174

Contents lists available at ScienceDirect

Economic Modelling

journal homepage: www.elsevier.com/locate/econmod

SIDD: An adaptable framework for analysing the distributional implications of policy alternatives where savings and employment decisions matter[☆]

Justin van de Ven[a,b]

[a] NIESR, 2 Dean Trench Street, London SW1P 3HE, UK
[b] MIAESR, The University of Melbourne, Australia

ARTICLE INFO

Keywords:
Dynamic programming
Savings
Labour supply
Microsimulation

ABSTRACT

The Simulator of Individual Dynamic Decisions, SIDD, is publicly available software for analysing the distributional effects of policy alternatives. SIDD is a framework, rather than a model, in the sense that it is designed to facilitate adaptation to alternative country and policy contexts. The microsimulation framework can generate panel data describing a wide range of characteristics at annual intervals for each adult in an evolving population cross-section. Structural methods are employed to project savings and employment decisions, making SIDD a suitable tool for exploring the incentive effects of policy alternatives, and how these vary across the population and over time. The framework is also a valuable test-bed for empirical analyses of alternative behavioural assumptions, especially those concerning preferences for risk. In an effort to support good policy design and empirical analysis of savings and labour supply behaviour, SIDD has been made free for download from www.simdynamics.org.

1. Introduction

Good policy design is a fiendishly difficult business due to the multiplicity, complexity, and inherent uncertainty of the considerations that are involved. One consideration that is often poorly understood is the variable impact that policy can have when considered over alternative time horizons. A welfare benefit may, for example, be interpreted as redistributing income between different members of a population when its incidence is observed at a particular point in time, and be interpreted as redistributing income across the life-course of individuals when considered over longer time horizons. Alternatively, a policy may have very different distributional implications when considered at alternative points in time, especially when behavioural responses are taken into consideration. Interest in understanding how policy influences individual circumstances over alternative time spans is an important motivation for the development of dynamic microsimulation models. This paper describes a framework that has been developed to limit the set-up costs associated with a dynamic microsimulation model in which savings and employment decisions are projected based on life-cycle theory; the Simulator of Individual Dynamic Decisions, SIDD. The full framework – including source code – is currently free for practitioners to download via the website: www.simdynamics.org.[1]

SIDD is a framework for developing structural dynamic microsimulation models. Working backward through this description, SIDD facilitates development of "models" that generate the logical implications of a set of stylised assumptions about how society functions. It is a "microsimulation", in the sense that each adult from a population cross-section is individually represented within the framework. SIDD is "dynamic" because it projects each adult through time, and it is "structural" because savings and employment decisions are projected based on a theoretical description of behaviour that is assumed to be independent of the policy environment. Finally, SIDD is a framework, because it is designed to facilitate adaptation to alternative country and policy contexts.

Like all dynamic microsimulation models, SIDD is fundamentally designed to explore the distributional effects of policy alternatives through time. Two terms referred to in the preceding sentence should be clarified. "Policy" is intended here to extend beyond explicit government decisions – concerning, for example, transfer payments – to include any aspect of

☆ SIDD is written in Fortran, with parallelisations implemented using OpenMP and MPI. Many thanks to Martin Weale, Guoda Cibaite, attendees at the European meeting of the International Microsimulation Association 2016 (Budapest), three anonymous reviewers, and the journal editor for useful comments on previous drafts. Program development received financial support from the Joseph Rowntree Foundation, HM Treasury, HM Revenue and Customs, the UK Department for Work and Pensions, the UK Institute for Arts and Actuaries, NAPF, AgeUK, the European Commission, the University of Melbourne, the Economic and Social Research Council, and the Leverhulme Trust. The usual disclaimer applies.
E-mail addresses: jvandeven@niesr.ac.uk.
[1] Public users of SIDD are requested to reference the current paper in any publication generated using the framework.

http://dx.doi.org/10.1016/j.econmod.2017.02.007
Received 30 August 2016; Received in revised form 31 January 2017; Accepted 12 February 2017
0264-9993/

the simulated environment. In context of SIDD, the policy environment encompasses a wide range of factors, including mortality rates, fertility rates, and risks of migration. Furthermore, "distributional effects" refer to variation of effects over any characteristic that can differ between simulated micro-units (adults). Commonly considered characteristics for distributional variation include income, wealth, age, birth year, relationship status, and the existence of dependent children. Whereas *static* microsimulation models are capable of exploring the distributional implications of policy alternatives at a point in time, the distinguishing feature of *dynamic* microsimulation models is that they can do the same through time. Intended uses of the framework are discussed at further length in Section 2.

Dynamic microsimulation models suitable for analysing the distributional implications of public policy have been growing in number and sophistication since the ground-breaking work of Orcutt (1957); see Li and O'Donoghue (2013) for a recent review citing 66 such models for 19 countries. Development of dynamic microsimulation models has benefitted from the increasing availability of detailed microdata, improvements in analytical methods, the advent of generic software packages (e.g. GENESIS, Edwards, 2010; LIAM, O'Donoghue et al., 2009; JAS-mine, Richiardi and Richardson, 2016), and a steady rise in computing power. Nevertheless, constructing this type of model remains both technically and computationally challenging, and current implementations consequently all impose non-trivial stylisations of one form or another.

One of the most important stylisations commonly adopted in the dynamic microsimulation literature concerns the projection of micro-unit behaviour. The importance of reflecting agent decision making increases with behavioural sensitivity to variation of interest (e.g. policy counterfactuals), and with the bearing that behaviour has on projected characteristics of interest (e.g. government budgets). Such considerations are exaggerated as the projected time-horizon is lengthened, due to feedback effects of behaviour on individual circumstances, and are therefore particularly relevant for dynamic microsimulation models that project circumstances well beyond a short (five year) time horizon. Nevertheless, fewer than one third of the models surveyed by Li and O'Donoghue (2013) are identified as using "behavioural equations" to project decisions through time.[2] Furthermore, even where behavioural variation is projected through time, it is common for these projections to either focus exclusively on employment, or use reduced form equations that are ill-suited to respond to evolving incentives; this is the case, for example, for all three of the dynamic microsimulation models for the UK cited by Li and O'Donoghue (2013) that include behavioural projections (PenSim2, SAGE, and a model produced at the IFS described in Brewer et al., 2007).[3]

Current best practice economic analysis uses the life-cycle framework to project behavioural responses to altered savings incentives. The approach assumes that savings and employment decisions are made as though to maximise expected lifetime utility. It provides an internally coherent basis for considering behavioural responses to policy alternatives, by assuming that the analytical description for utility is structurally independent of the policy context. Although some formulations of the life-cycle model imply analytically convenient closed forms (e.g. Pylkkänen, 2002), most academic attention focusses on specifications that are better adapted to capture the influence of uncertainty on decisions (e.g. Browning and Lusardi, 1996), and which require numerical Dynamic Programming (DP) methods to solve. DP methods are, however, computationally demanding (e.g. Rust, 2008) to an extent that necessitates use of bespoke software routines to reflect anything approaching a realistic policy context on prevailing computing technology. Existing generic software packages – including Matlab, GENESIS, LIAM, and JAS-mine – either do not currently include such routines, or impose a computational overhead that makes use of DP methods impractical.[4] As a result, the field of study has been limited almost exclusively to dedicated academic researchers. Furthermore, the associated literature omits much of the diversity of individual specific characteristics that is a conspicuous feature of the microsimulation literature more generally, or has abstracted from empirical identification to focus on the implications of theory (as is prevalent in the Agent Based literature; see, e.g. Richiardi, 2014, and Tran, 2016, for an example from the contemporary literature). SIDD is designed to fill this gap.

The SIDD framework has been made available to facilitate implementation of a structural dynamic microsimulation model that uses DP methods to project individual employment, savings and investment decisions for a new country or policy context, without the need to engage in complex programming. The ultimate objective is to allow the analyst to focus on identifying appropriate model parameters for reflecting a given economic context, which is itself a non-trivial task. To the best of our knowledge, SIDD is the closest that a dynamic programming framework currently comes to approximating the richness of individual circumstances that is accommodated in the wider dynamic microsimulation literature. Much of the heterogeneity that is commonly suppressed in focussed academic studies is important to policy makers, motivating inclusion in the framework. Furthermore, we are unaware of any other DP framework that is adapted to project the evolving population cross-section forward through time, which is necessary to consider a wide range of distributional implications of policy, including those concerning (relative) poverty, inequality, and the government budget. SIDD is also unique to the DP literature in that it is designed to project individual circumstances both *forward and backward* through time, where backward projections are made necessary by the objective to describe the life-course of older individuals in a reference population cross-section.

Section 2 provides an over-view of the framework, and discusses intended uses in context of the wider microsimulation literature. Sections 3 to 12 describe each model characteristic in turn, and Section 13 concludes. In keeping with the objective of developing a model fit for use by non-specialists, this text avoids use of technical terminology and detail wherever possible; see van de Ven (2016a) for more technical discussion.

2. Framework overview and intended uses

2.1. A brief overview

SIDD is designed to start from detailed micro-data describing the circumstances of a reference population cross-section, and to project data for each adult at annual intervals forward and backward through time. At the end of a simulation, the framework saves both the projected panel data for the simulated population, and a selected set of summary statistics for secondary analysis. Forward projections reflect the evolving population cross-section, so that the reference population is augmented to accommodate maturating children and international migratory flows.

[2] This omission of an explicit allowance for behaviour response is also a stylisation that is commonly employed in the wider empirical literature; see for example Kuang et al. (2011).

[3] All three of these models simulate employment transitions based on probabilities that vary over a range of characteristics, including demographics (e.g. age, sex, relationship status, dependent children), educational attainment, health status, and past work experience. SAGE and the IFS model summarise these probabilities in the form of logit regression equations, which can be derived from a trans-log utility function, and are sometimes therefore described as 'structural'. Nevertheless, these models are denoted 'reduced form' here, because none of them is designed to project labour responses to changes in transfer policy (the explanatory variables being exogenously defined). This means that the parameters of the logit regressions are likely to be functions of prevailing transfer policy. As a consequence, additional detail would be required to predict responses of employment probabilities to alternative policy environments.

[4] Initial development of SIDD made extensive use of pre-programmed routines provided in Matlab. However, use of Matlab was abandoned approximately seven years ago because associated computation times were found to be impractical.

The model unit is the family, defined as a single adult or partner couple and their dependent children (sometimes referred to as a 'benefit unit'). The framework can be used to consider endogenous decisions regarding consumption, labour supply of adults, and the portfolio allocation across a range of assets that include safe and risky liquid investments, personal pensions, and other tax advantaged savings accounts. The default assumption is that employment and savings decisions maximise expected lifetime utility. The set of characteristics that can be represented explicitly in the solution to the utility maximisation problem include: year of birth^{d,i}, age^{d,i}, relationship status^{i}, number and age of dependent children^{i}, student status^{i}, education^{i}, health status^{i}, carer status^{i}, migration status, potential full-time labour income, savings held in tax advantaged savings accounts, private pension eligibility, private pension wealth, timing of access to private pension wealth (retirement)^{d}, contributory state pensions^{d}, family wealth not otherwise defined, and time of death^{i}. Of the 17 characteristics listed here, four are assumed to evolve deterministically (indicated by a superscript '*d*'), and all others may evolve with some uncertainty. Furthermore, nine of the characteristics are simulated independently from the utility maximisation problem (indicated by a superscript '*i*'), and all others may be influenced – either directly or indirectly – by utility maximising decisions.

Fig. 1 provides a diagrammatic overview of how SIDD projects individual circumstances for the evolving population cross-section forward through time. Starting at the top of the chart, the framework begins by loading in a reference data-set describing all simulated characteristics for each adult in a reference population cross-section. These data are combined with a full set of random draws that provide the detail required to reflect the effects of uncertainty throughout life for every potential adult in the simulation. These two sets of data are the raw material used to project an evolving population cross-section through time.

The framework is structured around three classes of individual specific characteristics. *Preallocated characteristics* can be projected entirely separately from the remainder of the framework. These include the timing of death (which depends only on age and birth year), education (which depends only on age and birth year), and health status (which depends upon age, birth year, and education). SIDD projects data for the full life-course for each of these three characteristics, for all adults in the reference population cross-section. The framework then steps through the simulated time horizon one year at a time to project all other simulated characteristics.

Utility independent characteristics are entirely un-related to any decision that is projected to maximise expected lifetime utility in the framework. While the preallocated characteristics are also independent of utility, the key feature distinguishing utility independent characteristics is that they must be projected one year at a time, and can therefore not be preallocated as described above. Utility independent characteristics must be simulated one year at a time due to the approach taken to simulate marriages in forward projections, which assumes that relationships are formed between adults within the simulated sample. This means that the full sample of adults marrying in each year needs to be identified, before each is sorted into a newly-wed couple. Relationship status is projected with reference to each adult's age, birth year, education status, health status, existing relationship status and dependent children. As both fertility and carer status depend upon marital status, these variables are included in the set of utility independent characteristics in addition to relationship status. All utility independent characteristics are projected one year into the future before the framework projects any utility dependent variables for the same year.

As indicated by Fig. 1, up to six decisions can be projected to maximise expected lifetime utility. This is achieved via numerical routines that are designed to obtain approximations to the set of decisions that maximise expected lifetime utility, given any feasible combination of adult specific characteristics (see van de Ven, 2016a, Sections 2 and 14, for technical detail). After utility maximising decisions are projected, the *utility dependent characteristics* are projected forward one year. Projections for utility dependent characteristics depend upon the individual specific characteristics and utility maximising decisions in the prevailing year, and the utility independent characteristics projected forward one year. Having projected all simulated characteristics forward one year, the framework generates international migratory flows, with reference to age, education, marital status, dependent children, disposable income, and past migrant status. This process is repeated for each succeeding year, until projections are obtained for the entire simulated time horizon.

Each of the characteristics referred to in Fig. 1 is discussed at further length in the sections that follow, where the sections are ordered to approximate the simulation process as outlined in the figure. Before discussing simulation specifics, however, this section sets out the intended fields of application for the framework and best-practice methods of use.

2.2. Fields of application

SIDD is fundamentally designed to explore the distributional implications of public policy alternatives through time; see the review by Li and O'Donoghue (2013) for detailed discussion of the uses of dynamic microsimulation models more generally. The central feature that distinguishes SIDD from comparable dynamic microsimulation models is the structural approach used to project savings and employment decisions. Otherwise similar dynamic microsimulation models adopt behavioural assumptions that avoid the need to accommodate dynamic programming methods. This is predominantly achieved in the contemporary literature by limiting structural consideration of behaviour to employment, either in the form of projected retirement decisions (e.g. SESIM, DYNAMITE, and SADNAP models), or labour supply more generally (e.g. LIAM, LIAM2, and NEDYMAS models).

There are two key costs associated with using DP methods to accommodate savings and employment decisions in a dynamic microsimulation model. First, there is the developmental cost of implementing the DP method. A rough approximation is that, starting from scratch, including DP methods for savings, investment and employment decisions is approximately twice as difficult again, as programming the routines that actually project individual circumstances through time (thereby increasing developmental time by a factor of three). SIDD has been made publicly available specifically to mitigate this first cost.

Secondly, there is the computational burden associated with solving a DP problem in a realistic policy context. An appreciation of the scale of this issue might be gained by noting that a standard simulation currently run using SIDD requires 220 million utility maximisation problems to be numerically solved, the results of which must be stored *before* the task of projecting individual specific circumstances through time can begin. High performance computing hardware – currently in the form of personal workstations or computing clusters – is required to make this computational burden feasible on prevailing technology. Even then, however, the scale of the computing problem limits the plausible range of individual specific characteristics that SIDD can accommodate. As a practical comparison, relative to a standard simulation using SIDD, the MIDAS model for Belgium (e.g. Dekkers et al., 2015) generates just under four times as many characteristics in at least one year for a total population that is almost 25 times the size, in approximately the same time, on computing resources that are just over one third as powerful.[5]

SIDD is consequently well adapted to explore analytical contexts in

[5] Running a standard simulation of SIDD on a workstation with dual Intel Xeon E5-2670 processes and 96GB of RAM takes 2.5 hours to project a population backward 69 years and forward 30 years from a reference cross-section, generating 80 characteristics in at least one year for each of 90,000 simulated adults. In contrast, a standard simulation using MIDAS completes in around the same time, but generates data describing approximately 300 characteristics in at least one year during a 60 year projected time horizon for 2.2 million individuals, using a computer with a single Intel i5-750 processor and 16GB of RAM.

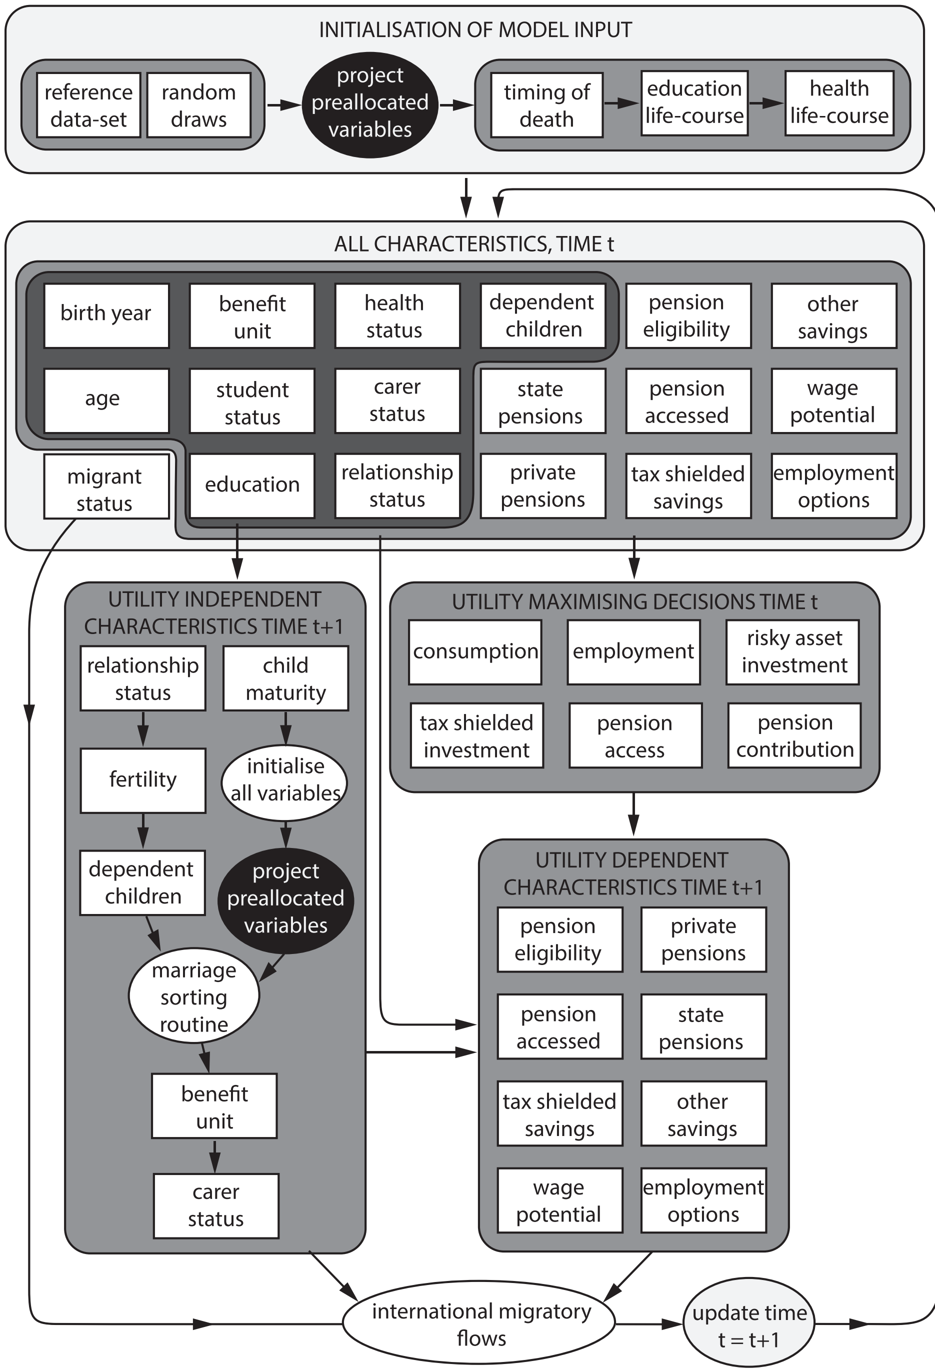

Fig. 1. Flowchart overview of how the Simulator of Individual Dynamic Decisions (SIDD) projects a population cross-section forward through time.

which the costs associated with the DP computational burden are more than offset by the coincident benefits of the structural approach to simulating savings and employment decisions; in other contexts, preferable modelling approaches exist. Two broad fields satisfy this condition. The first is evaluation of policy where behavioural responses are considered important. This is likely to be true of any policy counterfactual that is explicitly designed to influence behaviour, in which case SIDD can help to both guide expectations concerning the scale of behavioural responses, and to improve understanding of the incentives embodied by policy alternatives (which can be opaque). Behavioural responses to policy also become increasingly important, the further into the future interest extends. This is because behaviour today (e.g. savings) influences circumstances in the future (e.g. wealth), which can have a pronounced impact on behaviour (e.g. timing of retirement) in the future.

SIDD is currently used for public policy analysis by HM Treasury to explore the long-term implications of policy alternatives. Similarly, the framework was the principal tool used by the Joseph Rowntree Foundation to quantify the long-term implications for poverty and the government budget of the policy alternatives proposed in it's recently released anti-poverty strategy; see JRF (2016). Previous incarnations of the framework have also been used to explore a wide range of policy questions for the UK, with a particular focus on pensions; see e.g. Armstrong and van de Ven (2016); van de Ven (2013); Sefton and van de Ven (2009), and Sefton et al. (2008). The last two of these papers, for example, explore design of old-age safety-net payments in the UK. These papers highlight the militating effects associated with reducing the rate at which welfare benefits are withdrawn in respect of private income in retirement. As the rate of benefits withdrawal falls, price effects improve private incentives to save for retirement while income effects do the opposite. On balance, in a context where any adjustment in the overall cost of old-age welfare benefits is paid for by an adjustment in the general rate of income tax, the analysis reported in Sefton and van de Ven (2009) and Sefton et al. (2008) suggests that a withdrawal rate approximately mid-way between the extremes of no benefits withdrawal (a citizen's pension), and a 100% withdrawal rate (a minimum income guarantee) is likely to be preferred.

The second principal field for which SIDD is well adapted is empirical exploration of alternative assumptions concerning the life-cycle hypothesis. In this case, the DP framework is a central focus of interest. There is a growing literature that uses DP models of saving to test alternative behavioural hypotheses within the life-cycle framework, following the seminal study by Gourinchas and Parker (2002). SIDD is a useful framework for such work because it includes many of the characteristics that are likely to be important determinants of the two key decision margins of the domestic sector: consumption/savings, and labour/leisure. Furthermore, accommodation in SIDD of overlapping generations presents important technical advantages for exploring the empirical support for alternative savings assumptions; see van de Ven (2016b) for detailed discussion, and van de Ven and Weale (2010) for a practical example of econometric analysis based on the framework. The second of these two studies explores the empirical support for (time inconsistent) quasi-hyperbolic discounting, which has been suggested as a useful approach for capturing the influence of myopia within a rational-agent context. van de Ven and Weale (2010) presents weak econometric evidence in support of quasi-hyperbolic discounting, using differences in the liquidity of pension and non-pension savings for empirical identification.

2.3. Best-practice methods of analysis

2.3.1. Parameterisation and data

Although SIDD mitigates developmental costs associated with establishing a structural dynamic microsimulation model, re-specifying the framework to reflect a new economic environment remains a non-trivial task. Depending upon the desired functionality, there may be thousands of parameters that require specification, and a number of features have been adopted in the framework to facilitate this effort.

All model parameters are accessible through Excel spreadsheets alongside associated descriptive notes. The format of many of these parameters is designed to reflect statistics that are commonly reported by statistical authorities. For example, although the framework does not currently distinguish adults by their gender, the framework is based on mortality rates that do vary by gender as this is a common distinction made in official life-tables. Furthermore, aspects of the transfer system that are likely to require re-programming to reflect an alternative economic context have been compartmentalised from the remainder of the model code to facilitate their adaptation.

The vast majority of the parameters upon which SIDD depends are similar to those of any other dynamic microsimulation model. These parameters generally describe observable phenomena, such as mortality rates, fertility rates, marriage and divorce rates, wage parameters, and so on. The methods required to identify these parameters generally depend upon the available data: some parameters may be directly obtained from previously published estimates; some will require estimation using standard econometric techniques; and some parameters may require innovative statistical manipulation to overcome limitations of the available data.

The preference parameters upon which the structural projection of behaviour depends are unobservable, and must therefore be identified indirectly. As discussed in Section 2.2, a growing literature exists that uses econometric methods to estimate preference parameters for DP models of the type considered here. Such analysis is, however, very computationally demanding, and it may consequently be preferable to identify these parameters via manual calibration.

The SIDD framework that is available for download from the internet is supplied with a parameterisation for the UK. A detailed description of how this parameterisation was obtained is provided in van de Ven (2016b). Superseded versions of the framework have also been parameterised for Ireland (Callan et al., 2012) and Italy (European Commission Grant Number VS/2013/0208); see also van de Ven (2011) for empirical analysis of an earlier variant of the framework.

2.3.2. Policy analysis

The framework is best adapted to explore the effects of policy counterfactuals by taking differences between pairs of simulated projections, where the only variation between each simulation concerns the policy reform of interest. This approach provides a measure of how the policy reform would affect the population in a controlled context where all other features of the economy remained unaltered. From a policy maker's perspective, it focusses upon the influence of those features of the world that the policy maker can affect, ignoring features that are beyond their control. This is the approach taken in all of the examples of policy analysis previously conducted using SIDD, a sample of which is listed in Section 2.2.

Focussing on the effects of policy counterfactuals mitigates limitations associated with use of the framework as a forecasting tool. The most important limitation in this regard is that measurement of the full set of uncertainties associated with any future projection generated by SIDD is currently impractical. To explain this important point, it is useful to distinguish between three types of uncertainty associated with any projection. There are the uncertainties that are explicitly represented within a model's structure; in SIDD this covers a wide set of characteristics including the evolution of labour opportunities, investment returns, relationship status, and death. There are uncertainties associated with model parameters; in SIDD there is a large number of parameters governing everything from the likelihood of marriage and divorce to the influence of age on housing costs. And there are ‘other’ uncertainties, which include anything not explicitly represented in a model; in SIDD, these range from shifts in policy concerning same-sex marriages, through to the outbreak of war with a previously close trading partner. The influence of the first of these types of uncertainty can usually be reflected by a model; see Armstrong and van de Ven (2016) for a recent example using SIDD. The second type of uncertainty, associated with model parameters, is often more difficult to measure, as it can require a great deal of data to evaluate in any rigorous way. The third type of uncertainty – often the most important – is beyond the modelling scope.

The remainder of the paper describes the methods used to simulate each individual specific characteristic in turn.

3. Simulating health and mortality

As discussed in Section 2.1, health and mortality projections are simulated independently from the remainder of the framework. Each individual is allocated a health status at the time that they enter the simulated population, either recorded as part of the base data-set for adults in the reference population cross-section, or randomly allocated as part of the initialisation of characteristics for children maturing into the simulated sample frame. Furthermore, the framework base data include for each simulated adult a set of age-specific random draws from a uniform $[0, 1]$ distribution, which are referred to below as the ‘health vector’. The framework can distinguish the health status of dependent children, the health status of each simulated adult, associated carer responsibilities, and each simulated adult's time of death. As projections for health push computing technology to the limits of what is currently feasible, the framework has been designed only to project health and carer states forward, and not backward through time.

3.1. Mortality

The timing of death is the first characteristic that the framework simulates for each adult, in the set of 'preallocated variables' as displayed at the top of Fig. 1. The timing of mortality depends upon mortality rates that vary by age and simulation year, and which are commonly reported components of official life-tables. All adults in the reference cross-section are assumed to survive until at least the age at which they appear in the reference cross-section, with death possible in any subsequent year. Similarly maturing children are assumed to survive until their age of maturity. The age at death for each individual is evaluated by comparing the age specific elements of their health vector against the probability of death for an individual of the relevant age and year. If the element of the health vector is less than or equal in value to the respective probability, then the individual is assumed to die at the respective age.[6] Otherwise, they survive into the succeeding year. The age at death is consequently the first age at which the health vector is less than or equal to their respective probability of death.

3.2. Adult health status

Health status is the last of the three 'preallocated characteristics' that can be generated by the framework, and follows projection of mortality and education (discussed in Section 4). The framework is designed to distinguish between up to 10 discrete health conditions for each adult in each period projected forward through time. Behavioural solutions are structured around a health state described at the family level, $health_{i,a}^{j}$. In the case of single adults, this health state defines the health condition of the relevant adult. In the case of couples, the health state defines the 'health combination' of the two spouses. Simulated health conditions evolve through time, based on exogenously defined transition probabilities that vary by each adult's prevailing health condition, education, age, and year.[7]

The health state can influence families in a variety of ways. As noted above, the health condition of each adult can affect their likely health condition in the future. This feature is required to capture the persistence that is associated with many health conditions, which may have an important bearing on the life-course. The health condition of one adult in a couple can affect the carer responsibilities of their spouse (discussed in Section 5.5). The health state can also be defined to limit the discrete set of labour alternatives available to each adult, the probabilities of receiving a low wage offer and wages earned (Section 7), welfare benefits (Section 8), the likely evolution of relationship status in prospective years (Section 5.1), and non-discretionary costs, ndc^{h}. Family costs associated with adult health conditions can either be exogenously defined, or be set equal to the value of associated welfare benefits.

Care must be exercised when defining this aspect of the framework to ensure feasible computational times in context of prevailing technology. This can be done by limiting the number of discrete health states described by the framework. The framework is also designed to limit the computational problem by omitting health combinations associated with (near) zero probabilities, or those identified as of little interest by the analyst.

4. Education

Education is the second of the three preallocated characteristics generated for each adult by the framework (see Section 2.1). It is projected before health because health transition rates can vary by education status (as discussed in Section 3.2). The highest qualification held in all simulated years, $ed_{i,a}$, is allocated for each adult immediately after identifying the timing of death, and depends only upon age and education status in the reference data set, or birth year for children maturing into the simulated sample. It is possible to distinguish between up to five education states, one of which is reserved to reflect tertiary education, and the other four to reflect alternative levels of pre-tertiary education. The pre-tertiary education states can differ from one another in relation to the assumed probabilities of receiving a low wage offer and assumed wage premia (described in Section 7). In addition, individuals with tertiary education can be distinguished from non-tertiary educated in relation to the transition probabilities governing marriage and divorce (Section 5), and age specific evolution of latent wages (h in Section 7).

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. This approach for students is inverted in the projections backward through time. 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.

5. Relationship status, dependent children, and carers

As outlined in Section 2.1, after the framework finishes projecting the preallocated characteristics described above (mortality, education, and health) for all adults in the reference population cross-section, it generates data for the entire population one year at a time. The first set of characteristics that it generates when considering any given year are those that are independent of utility maximising decisions, but cannot be preallocated: relationship status, fertility, identification of family units, and carer status. This section describes simulation of these four characteristics in turn.

5.1. Relationship status

A 'relationship' is loosely defined in the framework as a cohabitating partnership, and may be parameterised to reflect alternative arrangements, including formal marriages and civil partnerships. In each period of a reference adult's life, their relationship status in the immediately succeeding period is uncertain, reflecting likelihoods of marriages, divorces, and widowhood. The transition probabilities that govern relationship transitions depend upon a reference adult's existing relationship status, age and birth year, the age and birth year of their spouse (if they have one), and may also vary with respect to their education status, health status, or the presence of dependent children. Probabilities of marriage and divorce are stored in a series of 'transition matrices', each cell of which refers to a discrete relationship/age/birth year combination; separate matrices are also stored that distinguish reference adults by education status, health status, and whether there are dependent children in the family. Similar transition matrices are used to model mortality (and widowhood), as described in Section 3.

When solving the lifetime decision problem, individuals are assumed to anticipate the probabilities of relationship formation, divorce, and death that govern relationship transitions. The decision problem is simplified by assuming that each individual expects to marry – if they marry – an identical clone of themselves. This

[6] Note that the health vector used to project mortality is the same as is used to evaluate the evolving health state, which serves to economise on the data that need to be recorded to permit replication of a simulation. Each element of the health vector is adjusted after it is referenced for identifying mortality to ensure that an independent random uniform number exists for projecting health state.

[7] See van Sonsbeek and Alblas (2012) for a dynamic microsimulation model designed specifically to explore disability related benefits.

assumption omits the uncertainty that would otherwise need to be accommodated concerning the characteristics of each potential spouse.

There are two principal approaches to projecting marital status in dynamic microsimulation models. The first assumes that the partners of newly married reference adults are drawn from outside the simulated population, and is used by SIDD for projections backward through time. In this case, the characteristics of spouses are simulated on the assumption that relationships form between identical clones, which has the advantage that it requires little additional data to be generated in respect of the spouse. The second approach adopts a 'closed model' specification that identifies married partners from within the simulated sample, and is used by SIDD to simulate the population forward through time. This difference in simulation approach between the forward and backward projections is motivated by the dual observations that a 'closed model' for relationships is both facilitated, and facilitates a model context that reflects the evolving population cross-section.

The cross-sectional data that are loaded into the framework to initialise the simulated population include the marital status of each represented adult. These data also include a personal reference number, and a family reference number. The same family reference number is recorded for each member of a cohabitating relationship. Matching of spouses in forward simulations is performed on a year-by-year basis. 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.

In backward simulations, the framework sets the reference adult of each family unit to the adult who is present in the reference cross-section (each of whom is identified in a separate family unit by assumption). In forward simulations, the use of a closed model approach for marriages requires selection of reference adults. In this case, the framework checks if one adult has accessed their pension wealth but the other has not (see Section 11.3), and if so, then the pension recipient is identified as the reference. Otherwise, the framework identifies the individual with the highest wage potential (see Section 7) as the family reference person.

At the time of a union in forward projections, jointly held assets are the sum of the assets held individually by each spouse. Widowhood in forward projections of the framework is based on the age of death simulated for each adult (described in Section 3.1). In backward projections, widowhood is identified randomly, based upon the mortality rates of the simulated adult (given the assumption of marriages between clones). In the event of widowhood, all assets and children of the family are assumed to reside with the surviving spouse. Divorce is simulated based on the transition probabilities applicable for the family reference person, in which case all assets and children are divided evenly between the respective spouses (to the nearest integer in the case of children).

5.2. Birth and aging of children

The framework is designed to take explicit account of the number and age of dependent children of reference adults. The birth of dependent children is assumed to be uncertain in the framework, and is described by transition probabilities that vary by the age, birth year, relationship status, and previously born children of a 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 family, children are assumed to remain dependants until an exogenously defined age of maturity. A child may, however, depart the modelled family prior to attaining maturity, if the reference adult experiences a relationship dissolution (to account for the influence of divorce).

Allowing for dependent children in the way set out in the preceding paragraph can lead to a very significant increase in the computational burden of the lifetime decision problem. If, for example, a family was considered to be able to have children at any age between 20 and 45, with no more than one birth in any year, and no more than six dependent children at any one time, then this would add an additional 334,622 states to the decision problem (with a proportional increase in the associated computation time).[8] In cases where children are not an issue of concern, the framework consequently allows associated uncertainty and heterogeneity to be suppressed. In this case, the number of dependent children in each family is described as a deterministic function of the age and relationship status of the reference adult. Where the number and age of dependent children are considered to be important, then the framework is made computationally feasible by limiting child birth to a fixed number of reference person ages. The framework may be directed, for example, to consider child birth only when the reference adult is aged 22, 26, or 33 years.

Capturing realistic family sizes in context of limited child birth ages will usually require that multiple births be allowed at each birth age. We might, for example, allow up to two children to be born at each of the three child birth ages referred to above, in which case the maximum number of children in a family at any one time would be limited to 6. In this example, the computation burden of the decision problem would be increased by a factor of 231, which is sufficiently constrained to make the solution to the decision problem feasible on contemporary computing technology.

Restricting the number of ages at which a child can be born in the framework raises a thorny problem regarding identification of the transition probabilities that are used to describe fertility risks. The framework calculates the required probabilities internally, based upon the assumed birth ages and fertility rates reported at annual intervals. This approach has been adopted both because statistical agencies tend to publish data at the annual age band level, and because it facilitates associated sensitivity analyses to be conducted around the number and precise birth ages assumed.

5.2.1. Child disability

The health status of children, $health^{c}_{i,a}$, distinguishes between two discrete alternatives that are designed to identify those with and without a persistent disability. In the case of children entering the population in the data for the reference cross-section, the cross-sectional data include a disability identifier. In the case of children entering due to simulated child birth, child health status is allocated by comparing the first element of the child's health vector with an exogenous incidence rate on the assumption that up to one disabled child may enter each family at each 'birth age'. The health status of children is assumed to remain unchanged until they mature into adulthood. Furthermore, child disability status is assumed to influence health and education status upon maturity (Section 5.4). It is consequently possible to use the framework to track the influence of disability throughout the life course, from birth through to death.

Child disability influences simulated welfare benefits (discussed in Section 8), the carer responsibilities of parents (discussed in Section 5.5), and family costs, ndc^{h}. Family costs associated with child disability can either be exogenously defined, or be set equal to the value of associated welfare benefits.

5.3. Projecting relationships and children into the past

As noted above, the framework assumes for backward projections that the spouses of reference adults enter and exit the simulated frame with their marriage. Backward projections of relationship status and fertility are complicated by the impact that relationship status is assumed to have on fertility, and the persistent effects that associated

[8] This assumes an age of maturity of 17.

changes have on the life-course. For example, the existence of a dependent child aged 10 years in a family unit reported in the reference data set has implications for the evolution of relationship status and the existence of siblings prior to the reference cross-section that are non-trivial to take into account. The framework uses a trial-and-error approach to address these complications, testing over alternative sets of random draws on the assumption that all individuals are single and childless at a defined minimum age (e.g. 16), until it a identifies a set that is consistent with each adult's relationship status and dependent children reported in the reference cross-section.

5.4. Child maturity

When dependent children reach an assumed age of maturity they depart their parental family and enter one of their own. In the backward projections, dependent children of reference adults enter the simulation frame when they are aged 17, and depart again in the year prior to their birth, in a similar fashion to the treatment of spouses of reference adults over marital transitions. In forward projections, dependent children can be followed into adulthood. This feature is implemented to reflect the evolving population cross-section, and is governed by two key model boundaries. The first is the maximum population size that the framework is directed to take into account, and the second is the number of periods into the future that projections are made. The more restrictive of these two boundaries determines the time horizon over which the evolving population cross-section is projected, and the framework reports this as part of its standard on-screen output. The time horizon of the projected population cross-section in turn determines whether a dependent child will mature and be simulated as an adult by SIDD; any child who matures beyond this cross-sectional time-horizon is dropped from the simulated sample.

When a child first achieves their maturity, a series of characteristics must be generated to continue their projection into adulthood. Each maturing child is assigned a unique person identifier. Their age and birth year are carried over from their parental family, and their year of entry into the simulated sample is recorded. All maturing children are identified as non-graduates. Education status of maturing children is otherwise allocated randomly, based on transition probabilities that can vary by child disability status, year and the education status of the reference adult in the parental family. Health status (discussed in Section 3) is randomly assigned, based on age and year specific transition rates. The health states to which maturing children are allocated can be limited in the case where a child is identified as disabled. All assets are set to zero for maturing children. This assumption is made because the framework does not account for child income, and unrequited transfers between families other than inheritances (see Section 8) are not included in the framework. Wage potential at age of maturity is based on a random draw from a log-normal distribution, the means and variances of which are age, year, and education specific.

The framework then proceeds to generate the complete life history of each maturing child for each of the three preallocated characteristics, as it does for each adult in the reference population cross-section (see Fig. 1, and Sections 3 and 4).

5.5. Carers

A “carer state” is generated for each family in each simulated period, $\text{carer}^{i,a}$, where carer families include one adult with carer responsibilities. The carer state evolves through time, based on exogenously defined transition probabilities that vary by the individual's prevailing carer state, the disability state of their spouse (see Section 3.2) and/or dependent children (see Section 5.2), age, and year. Carers can be limited to families with at least one adult who is sufficiently healthy, as defined by a pre-defined value of the health state.

Carers can be distinguished from other adults in regards to the benefits that they are eligible for (Section 8), their employment opportunities, and the time that they have available for leisure (Section 7).

6. The preference relation

Having evaluated the utility independent characteristics described in Section 5, the framework proceeds to project decisions that maximise expected lifetime utility (see middle pane of Fig. 1). In this case, behaviour is modelled as though a single ‘reference person’ makes decisions on behalf of all members of each family. Identification of reference people is described in Section 5.1.

The preference relation is a centrally important feature of the framework, and a wide range of alternative functional forms have been explored in the literature. The framework is consequently designed to facilitate experimentation of alternative assumptions,[9] and comes supplied with a nested CES utility function that is standard in the contemporary literature.

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

$$U_{i,a} = \frac{1}{1-\gamma}\left\{u\left(\frac{c_{i,a}}{\theta_{i,a}}, l_{i,a}\right)^{1-\gamma} + \Delta_{i,a} + + E\left[\beta_1 \sum_{j=a+1}^{A} \beta_0^{j-a}\left(\phi_{j-a,a}^{b} u\left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,j}\right)^{1-\gamma} + \left(1-\phi_{j-a,a}^{b}\right)\zeta_1(\zeta_0 + w_{i,j}^{+})^{1-\gamma}\right)\right]\right\} \quad (1)$$

$$u\left(\frac{c_{i,a}}{\theta_{i,a}}, l_{i,a}\right) = \left(\left(\frac{c_{i,a}}{\theta_{i,a}}\right)^{(1-1/\varepsilon)} + \alpha^{1/\varepsilon} l_{i,a}^{(1-1/\varepsilon)}\right)^{\frac{1}{1-1/\varepsilon}} \quad (2)$$

where $\gamma > 0$ is the (constant) coefficient of relative risk aversion; E is the expectations operator; A is the maximum age of survival; β_0 and β_1 are discount factors; $\phi_{j-a,a}^{b}$ is the probability of someone from birth year b living to age j, given survival to age a; $c_{i,a} \in R^+$ is discretionary composite non-durable consumption of the family of individual i at age a; $l_{i,a} \in [0, 1]$ is the proportion of family time spent in leisure; $\theta_{i,a} \in R^+$ is the family's equivalence scale; $\Delta_{i,a}$ represents the influence of decision costs on utility; the parameters ζ_0 and ζ_1 reflect the “warm-glow” model of bequests; and $w_{i,a}^{+} \in R^+$ is liquid net wealth when this is positive and zero otherwise. $\varepsilon > 0$ is the elasticity of substitution between equivalised consumption $(c_{i,a}/\theta_{i,a})$ and leisure $(l_{i,a})$ within each year. The constant $\alpha > 0$ is referred to as the utility price of leisure.

The labour supply decision is assumed to be made between discrete alternatives. No upper limit is imposed on the number of discrete alternatives, so that the labour decision can be made to approach a continuous margin.[10] Where adults are explicit, then a separate labour supply decision is allowed for each adult. Where health is explicit, then labour supply options can be constrained to reflect work-limiting conditions and carer responsibilities. Furthermore, employment options can be restricted in response to the lack of a simulated job offer to reflect involuntary unemployment (the likelihood of which can vary by age, education, health, relationship status, and potential full-time wage rate). An age of mandatory retirement is also included in the model parameters to limit the scale of the computational problem. We return to discuss labour supply in Section 7.

The equivalence scale θ is included in the preference relation to

[9] This is done by limiting the numerical search routines to those that do not depend upon first order conditions (which may not exist for some functions), and by compartmentalising calls to the utility function in a single program subroutine.

[10] The search routine used to identify utility maximising labour supply decisions searches over all feasible employment alternatives, implying that increasing employment options can substantively increase computation times.

reflect the important influence that family size has been found to have on the timing of consumption (e.g. Fernandez-Villaverde and Krueger, 2007). Similarly, decision costs are included in the preference relation to allow the framework to reflect behavioural rigidities that have been cited as important for understanding retirement savings decisions (e.g. Carroll et al., 2009). These costs are accommodated by reducing the value of Δ whenever behaviour deviates from pre-assigned default options in relation to private pensions and tax advantaged savings accounts.

The framework incorporates an allowance for behavioural myopia, through its assumption of quasi-hyperbolic preferences, following Laibson (1997). Such preferences are interesting because they are time inconsistent, giving rise to the potential for "conflict between the preferences of different intertemporal selves" (Diamond and Köszegi, 2003, p. 1840). Furthermore, the framework assumes that all discount parameters are the same for all individuals, and are time invariant. It also assumes that families are aware of any time inconsistency that their preferences display, a condition sometimes referred to as 'sophisticated myopia'. These assumptions rule out a number of interesting behavioural phenomena, including the capacity to reflect systematic population heterogeneity with respect to temporal biases (e.g. Gustman and Steinmeier, 2005), and procrastination (e.g. O'Donoghue and Rabin, 1999). Such effects could be accommodated without a qualitative increase in computational burden. Nevertheless, they are omitted here because the limited empirical analysis that we have conducted has failed to reveal important behavioural margins that such effects would help to explain. This is one principal research thread that we hope to pursue in the future.

The warm-glow model of bequests simplifies the utility maximisation problem, relative to alternatives that have been considered in the literature.[11] Including a bequest motive in the framework raises the natural counter-party question of who receives the legacies that are left. We return to this issue in Section 8.

7. Labour income dynamics

Labour income dynamics are primarily based upon 'potential full-time labour income', h, modelled at the family level:

$$log\left(\frac{h_{i,a}}{m_{i,a}}\right) = \psi_{i,a-1} log\left(\frac{h_{i,a-1}}{m_{i,a-1}}\right) + \kappa_{i,a-1}\frac{(1 - l_{i,a-1})}{(1 - l_{ft})} + \omega_{i,a-1} \tag{3a}$$

$$\begin{aligned} m_{i,a} &= m(n_{i,a}, ed_{i,a}, a, b) \\ \psi_{i,a} &= \psi(n_{i,a}) \\ \kappa_{i,a} &= \kappa(n_{i,a}, a) \\ \omega_{i,a} &\sim N(0, \sigma_\omega^2(n_{i,a}, ed_{i,a})) \end{aligned} \tag{3b}$$

where the parameters $m(.)$ account for wage growth, which in turn depend on relationship status $n_{i,a}$, education $ed_{i,a}$, age a, and birth year b. $\psi(.)$ accounts for time persistence in earnings, $\kappa(.)$ is the return to another period of experience, $(1 - l_{ft})$ is the leisure cost of full-time employment by all adult family members, and $\omega_{i,a}$ is an identically and independently distributed family specific disturbance term. The variance σ_ω^2 is defined as a function of relationship and education status. The only exception to Eq. (3a) is when a reference adult changes their education status (see Section 4). In this case, a new random draw is taken from a log-normal distribution, the mean and variance of which are specific to the family's age, birth year, relationship, and education.

Eq. (3a) is a parsimonious wage specification that has been explored at length in the literature (e.g. Sefton and van de Ven, 2004). The inclusion of an experience effect κ in the wage specification is what makes h dependent on utility maximising decisions. It also invalidates two-stage budgeting. Nevertheless, we have found this to be a useful feature in capturing rates of employment observed early in life, when labour income tends to be relatively low. The inclusion of an experience effect also ensures that any characteristic which does not feature explicitly in equation (3), but does influence employment decisions – such as health status, as described in Section 3 – will also affect the evolution of potential full-time labour income.

The employment income projected for a family in any year, $g_{i,a}$, depends upon the family's potential full-time labour income, h, and a series of adjustment factors that vary depending upon labour supply, previous access of pension wealth, education, and health. Each considered labour alternative is associated with a unique wage factor, which increases in labour participation, and is indexed to one when all adult family members are full-time employed. A separate factor allows labour income to respond systematically to past decisions to access pension wealth, which might reflect pre-conditions concerning contractual arrangements associated with pension access. Additional wage factors are used to reflect the existence of job offers, the wage premia associated with sub-tertiary levels of education, and the effects on wages of poor health.

8. The budget constraint

Eq. (1) is maximised, subject to an age specific credit constraint imposed on liquid net wealth, $w_{i,a} \geq D_a$ for the family of reference adult i at age a.[12] In context of income uncertainty, and a preference relation where marginal utility approaches infinity as consumption tends toward zero, rational individuals will never choose to take on debt equal to or greater than the discounted present value of the minimum potential future income stream that they face (however unlikely that stream might be). This rule is used to define D_a, subject to the additional constraint that all debts be repaid by age $a_D \leq A$.[13] Intertemporal variation of $w_{i,a}$ is, in most periods, described by the simple accounting identity:

$$w_{i,a} = w_{i,a-1} + \tau_{i,a-1} + ur^h_{i,a-1} - c_{i,a-1} - ndc^x_{i,a} + k_{i,a-1} + B_{i,a-1} \tag{4}$$

where τ denotes disposable income, ur^h is un-realised returns to owner-occupied housing, c is discretionary non-durable composite consumption, ndc^x is non-discretionary expenditure on category x, k represents net investment flows with tax advantaged savings accounts, and $B_{i,a-1}$ is the value of bequests received.

Non-discretionary costs are accommodated to reflect the possibility that some minimum expenditure is required to obtain the necessities of life (sometimes referred to as "committed expenditure"). Non-discretionary costs are disaggregated into child care, housing (rent and mortgage interest), health, and 'other' categories of expenditure to facilitate simulation of welfare benefits that make explicit reference to any one of these 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 family member. Non-discretionary health costs, ndc^h, can either be set equal to the value of dedicated welfare benefits generated in respect of recognised health conditions, or to the value of exogenously supplied model parameters. Non-discretionary housing expenditure is comprised of rent and mortgage payments, $ndc^{hg} = rent + mort$, and is described in Section 8.2. 'Other' non-discretionary expenditure, ndc^o, is defined in terms of equivalised (non-housing/non-child care/non-health) consumption, and can vary by age and year.

The only potential departures from Eq. (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.

[11] See, for example, Andreoni (1989) for details regarding the warm-glow model.

[12] Note that $w_{i,a}^+$ referred to above is related to $w_{i,a}$, with $w_{i,a}^+ = 0$ if $w_{i,a} < 0$, and $w_{i,a}^+ = w_{i,a}$ otherwise.

[13] The lower bound D_a is assumed to be the same for all households, as this helps to simplify the numerical search routines.

Wealth effects at the time a family accesses its pension wealth are discussed in Section 11. In relation to marital transitions, backward projections assume that spouses are identical clones (see Section 5.1 for discussion), so that wealth is halved in context of a dissolution and doubled in context of a formation. In forward projections, 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 family 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.

The methods used to simulate transfer policy and to project returns to liquid net wealth are now described, before addressing complications introduced when wealth is projected backward through time.

8.1. Simulated transfer policy

As the framework has been designed to undertake public policy analysis, particular care was taken concerning formulation of the module that simulates the effects of taxes and benefits. The framework allows the measures of income accruing to each adult family member to be accounted for separately, so that it can reflect taxation of individual incomes applied in many countries. The tax function assumed for the framework is represented by:

$$\tau_{i,a} = \tau\begin{pmatrix} b,\ a,\ n_{i,a},\ n_{i,a}^{c},\ health_{i,a}^{j},\ health_{i,a}^{c},\ care_{i,a}^{j}, \\ l_{i,a}^{j},\ g_{i,a}^{j},\ hh_{i,a},\ mh_{i,a},\ w_{i,a}^{h},\ rent_{i,a},\ mort_{i,a}, \\ rr_{i,a}^{h},\ w_{i,a}^{nh,j},\ r_{i,a}^{nh}w_{i,a}^{nh,j},\ pc_{i,a}^{c,j},\ pc_{i,a}^{nc,j},\ py_{i,a}^{j}, \\ k_{i,a}^{TA},\ w_{i,a}^{TA},\ r_{t}^{TA}w_{i,a}^{TA},\ ndc_{i,a}^{c,h},\ \eta_{b,a} \end{pmatrix} \tag{5}$$

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 $n_{i,a}^{c}$; health status of each adult j in the family, $health_{i,a}^{j}$; health status of each child, $health_{i,a}^{c}$; carer status of each adult, $care_{i,a}^{j}$; labour supply of each adult, $l_{i,a}^{j}$; the labour income of each adult, $g_{i,a}^{j}$; indicator variables for home-owners, $hh_{i,a}$, and mortgage holders, $mh_{i,a}$; the net owner-occupied housing wealth held by the family, $w_{i,a}^{h}$; the rent paid by non-home-owners, $rent_{i,a}$; the mortgage interest paid by mortgage holders, $mort_{i,a}$; the realised returns to (gross) housing wealth, rr^{h}; the non-housing net liquid wealth held by each adult in the family, $w_{i,a}^{nh,j}$; the investment return on liquid net wealth of each adult in the family, $r_{i,a}^{nh}w_{i,a}^{nh,j}$ (which may be negative); the concessional and non-concessional pension contributions made by each adult in the family, $pc_{i,a}^{(n)c,j}$; the (retirement) pension income received by each adult in the family, $py_{i,a}^{j}$; net contributions to tax advantaged savings accounts made during the prevailing year, $k_{i,a}^{TA}$ (which may be negative); the wealth held in tax advantaged savings accounts by the family, $w_{i,a}^{TA}$; the income earned on savings in tax advantaged accounts, $r_{t}^{TA}w_{i,a}^{TA}$; non-discretionary child care costs, $ndc_{i,a}^{c}$; non-discretionary health costs $ndc_{i,a}^{h}$; and a tax residual, $\eta_{b,a}$. The tax residual is designed to correct for differences between simulated and sample moments of disposable income, representing measurement error and (observable) departures between simulated tax and benefits policy and transfer policy as it was applied. All other inputs to the tax function are described in other sections of this paper.

Calculating taxes with respect to wealth held at the beginning of a period (as it is here) implies that disposable income is made independent of consumption in the same period. This is advantageous when consumption is a choice variable, as it implies that the numerical routines that search for utility maximising values of consumption do not need to evaluate disposable income for each consumption alternative that is tested.

8.2. Dis-aggregating liquid net wealth

Liquid net wealth includes all assets other than those that are otherwise explicitly represented in the framework. This composite asset is divided into three sub-categories by the framework: net wealth held in owner occupied housing $w_{i,a}^{h} \in [0, \infty)$; non-housing risky assets $w_{i,a}^{r} \in [0, \infty)$; and non-housing safe assets $w_{i,a}^{s} \in [D, \infty)$; $w_{i,a} = w_{i,a}^{h} + w_{i,a}^{s} + w_{i,a}^{r}$.

Given a measure of liquid net wealth, the framework begins by distinguishing housing from non-housing wealth ($w^{nh} = w^{s} + w^{r}$). Although formal modelling of housing investment decisions is analytically feasible, it is also computationally burdensome.[14] Computational feasibility of the framework is maintained by using reduced form equations to identify home owners (hh) and mortgage holders (mh), to project the value of net housing wealth (w^{h}) and mortgage debt (md^{h}), and to distinguish realised return on gross housing equity (rr^{h}) from unrealised returns (ur^{h}). Assumed interest rates are assumed to evaluate mortgage interest costs for mortgage holders ($mort$), and rental costs for non-home owners ($rent$) are based on simulated year, family size, and potential full-time labour income (h).

The portfolio allocation decision is represented in the framework as a choice concerning the proportion of non-housing liquid net wealth invested in risky assets during each year, $\rho_{i,a} \in [0, 1]$, and is only possible if non-housing liquid net wealth is positive. The division of non-housing liquid net wealth into safe and risky assets affects only the effective rate of return, $r_{i,a}^{nh} = \rho_{i,a}r_{t}^{r} + (1 - \rho_{i,a})r_{i,a}^{s}$, where r_t^r is the return to risky assets, and r^s the return to safe assets/debt. The rate of return to risky assets $\ln(r_{t}^{r}) \sim N\left(\mu_{r} - \frac{\sigma_{r}^{2}}{2}, \sigma_{r}^{2}\right)$ is assumed to be the same for all families at any point in time, t. The interest rate on safe liquid assets is assumed to depend upon whether $w_{i,a}^{s} = (1 - \rho_{i,a})w_{i,a}^{nh}$ indicates net investment assets, or net debts. Where $w_{i,a}^{s}$ is (weakly) positive, then r^s takes the value r^I. When $w_{i,a}^{s}$ is (strictly) negative then, r^s is designed to vary from r_l^D at low measures of debt to r_u^D when debt exceeds the value of working full time for one period (g^{ft}):

$$r^{s} = \begin{cases} r^{I} & \text{if } w^{s} \geq 0 \\ r_{l}^{D} + (r_{u}^{D} - r_{l}^{D})\min\left\{\frac{-w^{s}}{g^{ft}}, 1\right\},\ r_{l}^{D} < r_{u}^{D} & \text{if } w^{s} < 0 \end{cases} \tag{6}$$

Specifying $r_{l}^{D} < r_{u}^{D}$ reflects a so-called 'soft' credit constraint in which interest charges increase with loan size. The model parameters r^{I}, r_l^D, and r_u^D take fixed values when solving for utility maximising decisions, and are allowed to vary when simulating the intertemporal evolution of a population.

8.3. Backward projections and simulated inheritances

Two special complications are addressed by the framework when projecting liquid net wealth backward through time. The first is to ensure that decisions (especially consumption) projected for year $t - 1$ are consistent with the characteristics (especially liquid net wealth) in year $t - 1$, given the characteristics prevailing in year t. That is, if you suppose that consumption in year $t - 1$ was some value $\hat{c}$, then given liquid net wealth in period t we can use Eq. (4) to project back what liquid net wealth must have been in period $t - 1$, which might be $\hat{w}$. But the solution to the utility maximisation problem, given wealth $\hat{w}$, will imply a value of consumption at time $t - 1$, which might be $\bar{c}$. A simple search routine is used by the framework to ensure that $\bar{c} = \hat{c}$ in backward projections.

The second problem addressed by the model when projecting liquid net wealth backward through time concerns the assumed history of random events. As an individual ages, their assets portfolio will

[14] See, Aydilek (2013) for a dynamic programming model of housing investment; Creedy et al. (2015) explore policy experiments using a stylised two period representative agent model of savings and housing.

generally evolve in response to their accumulated life-history, responding to positive and negative shocks to a wide range of factors including labour market success, investment returns, health, relationship status, and so on. An individual who is in the top wealth decile at a given age is consequently likely to have experienced more favourable forms of variation during their lifetime than an otherwise similar individual in the bottom decile. Failure to accommodate this feature can result in unrealistic projections for wealth backward through time. For example, in cases where an individual is associated with insufficiently favourable variation in backward projections to reflect the assets they hold in the simulation reference period, the backward projections will indicate unrealistically high wealth holdings early in the adult lifetime.

The framework uses receipt of ‘inheritances’ as a tool in the backward projections for ensuring that random events that individuals are assumed to be subject to during their lives are consistent with the measure of liquid net wealth observed in the reference cross-section. Receipt of inheritances is only modelled in backward projections by the framework. The framework uses a search routine to allocate inheritances received so that the liquid net wealth of each adult in the reference cross-section falls below a threshold value at the beginning of their simulated lifetimes.

Implicit in the specification of preferences described by Eq. (1) is the assumption that inheritances may be left at the time of death of each adult. By definition, no reference adult described by the cross-sectional data from which model projections are made can have died prior to the year for which the cross-sectional data were observed. Hence, inheritances can only be left in periods projected forward through time. Where inheritances exist in forward projections, then these are assumed to be received by the surviving spouse, if one exists. Otherwise, inheritances are assumed to exit the frame of the simulation.

9. Tax advantaged savings accounts

Tax advantaged savings accounts in the framework are a hybrid retirement savings vehicle with three key features, based upon Individual Savings Accounts (ISAs) available in the UK. First, investment income and capital gains within a tax advantaged account are tax free, both at the time earned and upon withdrawal. Secondly, annual contributions are subject to upper limits. And thirdly, no time limit is imposed on when accumulated funds can be withdrawn. The first of these elements encourages contributions into the scheme, the second discourages withdrawals, while the third relaxes the liquidity disincentives associated with traditional pension schemes.

Each family is assumed to be able to contribute to a single tax advantaged account. Annual contributions to the tax advantaged account are made out of post-tax income, and are subject to a per-period cap that doubles where the family is comprised of an adult couple. Although a distinction currently exists in the UK between so-called ‘cash’ and ‘stocks-and-shares’ ISAs, the framework is adapted to consider only one of these types at a time. At the start of each period, all wealth held in an tax advantaged account is assumed to accrue the same rate of return, r_t^{TA}, which can be specified as uncertain. Uncertain returns to a tax advantaged account are assumed to be perfectly correlated with the returns to the risky liquid asset (r_t^r in Section 8.2). In most periods, wealth held in a tax advantaged account, w^{TA}, is assumed to vary intertemporally as described by the equation:

$$w_{i,a}^{TA} = r_{t-1}^{TA}\, w_{i,a-1}^{TA} + k_{i,a}^{TA}$$
$$ln(r_t^{TA}) \sim N\left(\mu_{TA} - \frac{\sigma_{TA}^2}{2}, \sigma_{TA}^2\right)$$
$$corr(r_t^{TA}, r_t^r) = 1 \tag{7}$$

where $k_{i,a}^{TA}$ denotes net contributions into the scheme (negative when there are net out-flows), and $corr(.)$ denotes the correlation coefficient. The only departure from Eq. (7) is when the relationship status of a reference adult is identified as changing, in which case associated fluctuations in tax advantaged accounts are modelled in the same fashion as described for liquid net wealth (see Section 8).

As noted in Section 6, the preference relation assumed for analysis also allows for the possibility that contributions to tax advantaged accounts are influenced by decision costs, $\Delta_{i,a}^{TA}$. In this case, utility is assumed to decline discontinuously when the first contribution to a family's tax advantaged account is made.

10. Contributory state pensions

The contributory state pension in the framework is a flat-rate benefit, based upon the UK basic State Pension, rights to which are accrued through accreditation in respect of contributions during the working lifetime. The framework tracks the number of years, $yr_{i,a}^{CSP}$, for which each family, i, at age a, has been accredited with contributions, up to the maximum defined by the number of years required for a full contributory state pension for each adult family member. Accreditation for contributions is derived if the earnings of an adult exceed a minimum threshold, and can also be allowed for in respect of child care (non-employment during peak child-rearing ages), or involuntary unemployment (periods in which a low-wage offer is received – see Section 6). In most years prior to state pension age, the number of years of accreditation for contributions is defined by:

$$yr_{i,a}^{CSP} = yr_{i,a-1}^{CSP} + k_{i,a-1}^{CSP} \tag{8}$$

where $k_{i,a-1}^{CSP}$ are the additional contributions accredited to family i at age $a-1$. The only exception to equation (8) is when the relationship status of a reference adult is identified as changing, in which case associated fluctuations in state pension rights are modelled in a similar fashion as described for liquid net wealth (see Section 8).

Each family is assumed to draw its contributory state pension from state pension age, a^{SPA}, which is permitted to vary between birth cohorts, and this public transfer is added to pension income for tax purposes. The value of the state pension payable to each family depends upon the contributions history of the family, the value of the full state pension assumed for the reference year, a growth rate applied until the time when the reference adult of the family attains state pension age, and another growth rate applied from state pension age. Two values of the full contributory state pension are taken into consideration; one for single adults, and another for adult couples. The framework assumes that each family is paid the greater of the single allowance, paid in respect of the number of complete contribution histories accrued by all adult family members, and the couple allowance, paid in respect of a single adult's contribution history for couples. The framework does not track each adult's contribution history separately, but instead assumes that all contribution years accrue to the reference adult up to the number of years required for a full contributory state pension, and to the spouse (if one exists) thereafter.

11. Private pensions

The description of private pensions in the framework is designed to capture the broad features of the contemporary UK pensions environment, which is conspicuously diverse viewed from an international perspective. The model distinguishes between Occupational Pensions (OP) that are conceptually run by companies on behalf of their employees, and up to five Personal Pension schemes (PP) that individuals provide for themselves. These schemes can differ from one another concerning the terms of pension contributions, and returns to pension wealth.

All 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. Where OP and PPs are run in parallel, then any family with a labour income in excess of a lower bound is assumed to participate in the OP, while any family in which the highest adult earner has a labour income within an income band can be given the option to participate in a PP. The income thresholds used to manage eligibility to private pensions can (but do not have to) overlap. Where multiple PPs are accommodated in the framework, then each family is considered to be eligible for a single PP in each simulated period, where the evolution of pension eligibility is determined by an auto-regressive random process.

11.1. Private pension contributions

Contributions to private pensions are defined as rates of employment income (implying that they are limited to families that work), and are distinguished by whether they are made by the employer or the employee. Employer contributions are assumed to be exempt from taxation, and labour income is reported net of these contributions. Employee contributions can be subject to taxation, and labour income is reported gross of these contributions.

Contributions to an OP are simulated exogenously in the framework. The framework assumes that any family that has not previously accessed its pension wealth, and which has labour income in the prevailing period in excess of a lower threshold, will contribute to the OP. The employee contribution rate to the OP, ${\pi_{ee}}^{OP}$, and employer contributions rate, ${\pi_{er}}^{OP}$, are defined as fixed proportions of family labour income. Hence, for any family that contributes to an OP in a given period, the value of the contribution is defined as: $pc_{i,a}^{OP} = (\pi_{ee}^{OP} + \pi_{er}^{OP})g_{i,a}$.

Each PP scheme, p, has a single employer contribution rate, $\pi_{er}^{PP,p}$, and a (minimum) employee contribution rate, $\pi_{ee,0}^{PP,p}$, which apply to labour income within an income band defined by upper and lower bounds. Contributions to a PP can either be imposed, or be simulated to maximise expected lifetime utility. Where multiple PPs are considered, then utility maximising pension contribution decisions are limited to whether or not to contribute (the extensive margin). Where only one PP is considered for analysis, then the family can choose both whether to make fresh contributions to their eligible scheme, and how much to contribute (the intensive margin) in excess of the imposed minimum, $\pi_{ee,i,a}^{PP,p} \geq \pi_{ee,0}^{PP,p}$.

When contributions to a PP maximise utility, then it is also possible to allow for the influence of assumed ‘default options’, which can vary between alternative PP schemes. This is achieved by allowing for ‘decision costs’ that discontinuously reduce welfare when a family decision departs from the pre-assigned default (represented by Δ in Eq. (1)). Defaults may be considered over both contribution rates and/or participation alternatives. Defaults concerning participation are defined when a family first becomes eligible to a given scheme (i.e. auto-enrolment or active opt-in), and thereafter track a family's pension decision in the preceding period.

11.1.1. Pension contribution caps

Aggregate private pension contributions (from both OPs and PPs) can be subject to three separate caps. Two of these caps are year specific, can be altered over three age intervals, and can be administered at either the individual or family level. The first of these period-specific caps defines the upper limit on pension contributions that are eligible for tax relief. Employer contributions are accommodated first within this cap, after which private contributions are included. Private pension contributions that can be accommodated within this first cap are referred to in the framework as *concessional contributions*, pc^c. Any contribution in excess of the first cap can also be subject to a second cap defining the maximum contributions that are permissable each simulated period. Any employer contributions not accommodated within the concessional contributions cap are preferentially accommodated within the second cap. If employer contributions are in excess of the sum of these two period-specific caps, then the excess of employer contributions is considered to be returned to the employer and no private contributions are permitted. Any private contribution accommodated within the second cap is recognised as a *non-concessional contribution*, pc^{nc}. The third cap on pensions is an upper bound on the maximum size of the private pension pot; any contribution made that would result in this third cap being breached is assumed to be taxed at the rate of 100%. The first two of the caps defined above can be allowed to vary through time, at year specific rates within a closed and bounded period, and at fixed rates beyond this period. The upper bound on the size of the pension pot is assumed to remain fixed in real terms through time.

11.2. Evolution of private pension rights during accrual

Where multiple PP schemes are considered for analysis, balances of PP wealth are assumed to be perfectly portable between schemes, and each family is assumed to hold all of their personal pension wealth in the scheme for which they are eligible at the given point in time. The same basic rules are used to project through time wealth held in occupational and personal pensions, $p = \{OP, PP\}$. Until the year in which a family accesses its pension wealth, returns to private pension savings attract a rate of return that can be defined as uncertain. When returns are assumed to be uncertain, then they are perfectly correlated with the returns to the risky liquid asset (${r_t}^r$ in Section 8.2). Hence, accrued pension rights do not hedge against uncertainty in the liquid asset portfolio. Intertemporal accrual of private pension wealth, w^p, is described by Eq. (9):

$$
\begin{aligned}
&w_{i,a}^{p} = max\left\{0,\ min\left[w^{p,max},\ r_{t-1}^{p}\ w_{i,a-1}^{p} + pc_{i,a}^{p}\right]\right\}\\
&ln({r_t}^{p}) \sim N\left(\mu_p - \frac{\sigma_p^2}{2},\ \sigma_p^2\right)\\
&corr(r_t^{p},\ r_t^{r}) = 1
\end{aligned}
\tag{9}
$$

where p distinguishes OP, and each of the PP schemes considered for analysis, $w^{p,\max}$ defines the maximum size of a (PP or OP) pension pot, and $corr(.)$ is the correlation coefficient. The returns accruing to alternative PP schemes can be allowed to vary to reflect, for example, differences in assumed management charges. Eq. (9)) 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.

11.3. Accessing pension wealth and retirement

The age at which pension dispersals are first accessed, a^P, can be either imposed or simulated to maximise expected lifetime utility. If both OP and PP schemes are modelled, then rights to both pensions are assumed to be accessed at the same time. At the time that pension wealth is accessed, a fixed fraction of accrued pension wealth (that may differ between occupational and personal pensions) is received as a lump-sum cash payment, and the remainder converted into an annuity. The framework can be used to simulate either fixed term or life annuities. Annuity rates for life annuities are calculated to reflect birth cohort-specific survival probabilities in the framework, subject to assumed rates of investment returns, real growth, and transaction costs levied at time of purchase. The tax treatment of both the lump-sum and pension annuity can also be specified.

When the timing of pension dispersals maximises lifetime utility, then the decision can be made subject to minimum thresholds on age and annuity income. Furthermore, limits can be imposed on a family's pension contributions and employment opportunities following pension take-up. Employment opportunities can be subject to both hard limits (on the ability to find employment following pension take-up), and soft limits (in the form of wage penalties imposed on pensioner families, discussed in Section 7).

11.4. Simulating pensions backward through time

Simulating pensions backward through time shares many similarities with the simulation of liquid net wealth described in Section 8. The principal innovations associated with the simulation of pensions concern the treatment of the timing of pension access when this is endogenous. Consider, for example, the case where the framework is projecting backward through time the circumstances of a family that is identified as having started to draw down their pension wealth in some period prior to the reference cross-section. The framework must then identify when the family first started to access its pension wealth. Ensuring incentive compatibility of backward simulations is complicated in this context by the existence of multiple ages at which pension access will be optimal, given the assumption that pension income was not taken up at some preceding age. SIDD includes an optimised search routine to address this complication, which tests over the full range of feasible alternatives (given circumstances described by the reference data-set).

12. International migration

As discussed in Section 2.1, international migratory flows are the last subject that the framework addresses before it proceeds to consider a prospective simulation year. These flows are included to permit the framework to reflect the evolving population cross-section. As the review by O'Donoghue et al. (2010) makes clear, there are a wide range of alternative approaches used to simulate the effects of migration in the microsimulation literature. Key modelling decisions include whether to model net migration or immigration and emigration separately, the variables that describe the likelihood of emigration, the approach taken to generate the characteristics of immigrants, and whether to accommodate re-entry of emigrants. These decisions depend upon the reasons for the respective model's development, and the data that are available for parameterisation.

Migration has been accommodated in SIDD to meet two key objectives. First, the framework should be capable of reflecting official projections for the age distribution of the population through time. Secondly, the framework should reflect the bearing that contemporary trends concerning migration would have on the distribution of income if they were to continue into the projected future. Although the first of these two objectives could be achieved by modelling net migration only, this approach would complicate achieving the second objective. SIDD is consequently designed to accommodate explicitly both immigration and emigration in each simulated period.

There are two principal approaches for generating the characteristics of recent immigrants in a microsimulation context (e.g. Duleep and Dowhan, 2008). The 'regression' based approach involves estimating a system of equations describing all of the characteristics of interest, and uses these equations to generate characteristics for new immigrants. Valid implementation of this approach is, however, exceptionally challenging in any context where more than a few characteristics are involved, as is the case in the current context. SIDD consequently generates the characteristics of new immigrants using the alternative approach, by 'cloning' families from 'donors' drawn from targeted population subgroups. This means that the approaches taken to simulate immigrants and emigrants share close similarities with one-another.

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 framework 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. It is also possible to distinguish between up to two source regions (in addition to natives) for migrant flows, *region*, so that the framework is capable of reflecting, for example, differences between migrants by whether their country of origin was the host country, the EU, or some other country.

13. Conclusions

This paper describes the Simulator of Individual Dynamic Decisions (SIDD), a framework designed to facilitate the development of structural dynamic microsimulation models. SIDD is fundamentally designed to explore the distributional implications of public policy alternatives through time. There has been a proliferation of models of this sort during the last few decades, supported by improvements in computing power, analytical methods, and data availability. These advances have also motivated release of a number of generic software tools to aid model development, including LIAM(2), GENESIS, and JAS-mine. The innovation of SIDD, relative to these existing generic software packages, is that it comes pre-programmed with the functionality required to simulate an evolving population cross-section, where savings and employment decisions are projected based on the life-cycle framework.

SIDD is designed to simulate the evolving circumstances of adults in a representative population cross-section forward and backward through time, and to project the evolving population cross-section forward through time. The framework is designed to allow for differences between adults regarding their year of birth, age, relationship status, number and age of dependent children, student status, education status, health status, carer status, migration status, employment status, labour income, savings in tax advantaged accounts, private pension eligibility, private pension wealth, timing of access to private pension wealth, state pension rights, family wealth not otherwise defined, and time of death. Decisions that can be simulated based on the life-cycle framework include (non-durable) discretionary consumption, labour supply, investments in risky assets, investments in tax advantaged savings accounts, private pension contributions, and the timing of access to pension wealth. Uncertainty can be taken into account with respect to prospective labour market opportunities, investment returns, education status, relationship status, dependent children, health status, and time of death. Particular care has been taken to allow the framework to reflect a detailed description of tax and benefits policy.

As alluded to above, SIDD is designed specifically to facilitate development of models for analysing the distributional implications of policy alternatives through time in contexts where behavioural responses are considered important. This is likely to be true wherever policy is explicitly designed to affect behaviour, in which case a model based on SIDD can help to both quantify anticipated responses and improve understanding of the incentives underling policy alternatives. The behavioural focus of SIDD is also increasingly important, the further into the future analytical interest extends, due to the effects that policy can have on individual circumstances, which may feed through to the decisions people make.

Furthermore, the fact that SIDD embeds a life-cycle structure within a realistic policy context makes it a useful test-bed for empirical evaluation of alternative behavioural assumptions. There remains extensive debate concerning the features that a preference relation should have, exaggerated in the decade following the financial crisis. In this regard, the fact that SIDD comes pre-packaged with a standard preference relation should be interpreted as an invitation for associated experimentation.

The SIDD framework – including underlying source code – has been made publicly available in recognition of the non-trivial programming hurdles associated with establishing a dynamic microsimulation model that uses dynamic programming methods to project savings and employment decisions. Until now, such models have been developed independently by specialist researchers, each of whom maintained their own model framework. We have been one of the participants in that literature. It is hoped that reducing the costs of entry to this interesting field of study will help to support associated research, encourage sharing of best-practice analytical methods, and ultimately improve the evidence base for good policy reform.

References

Andreoni, J., 1989. Giving with impure altruism: applications to charity and Ricardian equivalence. J. Political Econ. 97, 1447–1458.

Armstrong, A., van de Ven, J., 2016. The impact of possible migration scenarios after 'Brexit' on the state pension system. Economies 4, 1–13.

Aydilek, A., 2013. Habit formation and housing over the life cycle. Econ. Model. 33, 858–866.

Brewer, M., Browne, J., Emmerson, C., Goodman, A., Muriel, A., Tetlow, G., 2007. Pensioner poverty over the next decade: what role for tax and benefit reform?, IFS Commentary No. 103.

Browning, M., Lusardi, A., 1996. Household saving: micro theories and macro facts. J. Econ. Lit. 34, 1797–1855.

Callan, T., Keane, C., van de Ven, J., 2012. A framework for pension policy analysis in Ireland: PENMOD, a dynamic simulationmodel. In: Callan, T., O'Connell, J. (Eds.), Analysing Pensions: Modelling and Policy Issues. ESRI, Dublin.

Carroll, G.D., Choi, J.J., Laibson, D., Madrian, B.C., Metrick, A., 2009. Optimal defaults and active decisions. Q. J. Econ. 124, 1639–1674.

Creedy, J., Gemmell, N., Scobie, G., 2015. Pensions, savings and housing: a life-cycle framework with policy simulations. Econ. Model. 46, 346–357.

Dekkers, G., Desmet, R., Fasquelle, N., Weemaes, S., 2015. The social and budgetary impacts of the recent social security reform in Belgium. In: Salagean, I., Lomos, C., Hartung, A. (Eds.), The Yound and the Elderly at Risk. Intersentia, Cambridge.

Diamond, P., Köszegi, B., 2003. Quasi-hyperbolic discounting and retirement. J. Public Econ. 87, 1839–1872.

Duleep, H., Dowhan, D., 2008. Adding immigrants to microsimulation models. Social. Secur. Bull. 68.

Edwards, S., 2010. Techniques for managing changes to existing simulation models. Int. J. Microsimulation 3, 80–89.

Fernandez-Villaverde, J., Krueger, D., 2007. Consumption over the life cycle: facts from consumer expenditure survey data. Rev. Econ. Stat. 89, 552–565.

Gourinchas, P., Parker, J., 2002. Consumption over the life cycle. Econometrica 70, 47–89.

Gustman, A.L., Steinmeier, T.L., 2005. The social security early entitlement age in a structural model of retirement and wealth. J. Public Econ. 89, 441–464.

JRF, UK Poverty: Causes, Costs and Solutions. Joseph Rowntree Foundation, York, 2016.

Kuang, Y., Englebrecht, T., Gilley, O.W., 2011. A distributional analysis of the FairTax Plan: annual and lifetime income considerations. South. Econ. J. 78, 358–381.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Q. J. Econ. 112, 443–477.

Li, J., O'Donoghue, C., 2013. A survey of dynamic microsimulation models: uses, model structure and methodology. Int. J. Microsimulation 6, 3–55.

O'Donoghue, T., Rabin, M., 1999. Procrastination in preparing for retirement. In: Aaron, H.J. (Ed.), Behavioral Dimensions of Retirement Economics. Brookings Institute, Washington, 125–156.

O'Donoghue, C., Lennon, J., Hynes, S., 2009. The Life-cycle Income Analysis Model (LIAM). Int. J. Microsimulation 2, 16–31.

O'Donoghue, C., Redway, H., Lennon, J., 2010. Simulating migration in the PENSIM2 dynamic microsimulation model. Int. J. Microsimulation 3, 65–79.

Orcutt, G., 1957. A new type of socio-economic system. Rev. Econ. Stat. 58, 773–797.

Pylkkänen, E., 2002. Modelling wages and hours of work. In: Pylkkänen, E. (Ed.), Studies on Household Labor Supply and Home Production. Goteborgs Universitet, Gothenburg.

Richiardi, M., Richardson, R.E., 2016. JAS-mine: a new platform for microsimulation and agent-based modelling. Nuffield College Economic Working Papers 2016-W04.

Richiardi, M., 2014. The missing link: AB models and dynamic microsimulation. In: Leitner, S., Wall, F. (Eds.), Artificial Economics and Self Organization: Agent-Based Approaches to Economics and Social Systems. Springer International Publishing, Cham.

Rust, J., 2008. Dynamic programming. In: Durlauf, S.N., Blume, L.E. (Eds.), New Palgrave Dictionary of Economics. Palgrave Macmillan, New York.

Sefton, J., van de Ven, J., 2004. Simulating household savings and labour supply: an application of dynamic programming. Natl. Inst. Econ. Rev. 188, 56–72.

Sefton, J., van de Ven, J., 2009. Optimal design of means tested retirement benefits. Econ. J. 119, F461–F481.

Sefton, J., van de Ven, J., Weale, M., 2008. Means testing retirement benefits: fostering equity or discouraging savings? Econ. J. 118, 556–590.

Tran, C., 2016. Fiscal policy as a temptation control device: savings subsidy and social security. Econ. Model. 55, 254–268.

van de Ven, J., Weale, M., 2010. An empirical investigation of quasi-hyperbolic discounting. National Institute of Economic and Social Research Discussion Paper 355.

van de Ven, J., 2011. A structural dynamic microsimulation model of household savings and labour supply. Econ. Model. 28, 2054–2070.

van de Ven, J., 2013. Analysis of pension reform scenarios in a rational world: an application of the NIBAX behavioural micro-simulation model. Dep. Work Pension-. Work. Pap. Ser. 117.

van de Ven, J., 2016a. LINDA: a dynamic microsimulation model for analysing policy effects on the evolving population cross-section. NIESR Discussion Paper 459.

van de Ven, J., 2016b. Parameterising a detailed dynamic programming model of savings and labour supply using cross-sectional data. Int. J. Microsimulation.

van Sonsbeek, J.-M., Alblas, R., 2012. Disability benefit microsimulation in the Netherlands. Econ. Model. 29, 700–715.