Final Report

The feasibility of building an add-on component for the CAPITA model to analyse the impact on labour supply of income support policies

Nicolas Hérault, Guyonne Kalb and Justin van de Ven

Melbourne Institute of Applied Economic and Social Research

Acknowledgements

This research was commissioned by the former Australian Government Department of Education, Employment and Workplace Relations (DEEWR) under the Social Policy Research Services Agreement (2013–14) with the Melbourne Institute of Applied Economic and Social Research. The views expressed in this report are those of the authors alone and do not represent those of the department.

October 2014

Table of Contents

Summary

1. The objective of this report is to improve the evidence base for contemporary decisions concerning model development within the Department of Employment (hereafter the Department). An important back-drop to this report is current work being conducted jointly by the Treasury, the Department of Social Services, and the Department of Employment to develop a microsimulation model, named CAPITA. The report is motivated by the observation that CAPITA is currently not planned to project labour supply responses to policy alternatives that are a principal focus of interest to the Department.
2. Treasury has pre-existing experience in the development of microsimulation models that are capable of projecting labour supply responses. However, existing models available to Commonwealth departments (including the Department of Employment) that allow for labour supply responses to policy change – MITTS and STINMOD-B – do not reflect some aspects which have become relevant due to the evolving policy environment. Important examples include childcare, health, working late in life, and activity requirements.
3. This report sets out the alternative methods that represent current best practice in the analysis of labour supply responses to tax and transfer policies, including discussion of their respective advantages and disadvantages. Drawing on this field of research, the report provides clear recommendations for model development – including sensible terms of reference – based on evolving analytical capacity within the Department, and associated policy needs.
4. We outline the most practical and suitable approach (given the Department's interests) to extend CAPITA to permit consideration of labour supply responses to the policy environment, through the development of an add-on behavioural module. The report sets out sensible terms of reference, explores viable modelling alternatives in context of the contemporary literature, and considers the issues involved in practical implementation. Much of the detail in this report has been written with the objective of providing most of the technical detail required to issue a public tender for developing the module, or a road-map for in-house development and implementation.

Labour Supply Modelling Methodology

5. A range of current best-practice methods currently exists, with no single best approach. There are two key criteria over which existing methods vary: the description assumed for individual behaviour, and the dimensionality of the considered population.

Behavioural assumptions

6. Behavioural assumptions range from exogenously imposed employment outcomes at one extreme to formal models of decision making at the other extreme. We identify three broad classes of modelling approach:
   a. Studies that exogenously impose labour supply responses typically make broad generalisations that are easy to document but they do not provide any predictive information concerning how labour supply will respond to alternative policy contexts.
   b. Reduced-form regressions are designed to summarise in an accessible form correlations between observable characteristics that are described by survey data. This methodology summarises historical behaviour in a simple and transparent way. However, the parameters of the equations that are considered for analysis may themselves depend upon factors that are determined by policy makers. Hence, this methodology is better suited for exploring behavioural responses to historical policy reforms than for considering likely behavioural responses to policy counterfactuals.
   c. Structural models are founded on a theoretical framework that is assumed to be structurally stable, in the sense that it remains invariant to the policy environment. Assumptions concerning how decisions are made, combined with data on past behaviour, permit the model to predict behaviour in the future. The key advantage of a structural model is that it provides a coherent basis for exploring behavioural responses to policy counterfactuals. The central drawback of the approach is that the theoretical framework upon which the analysis depends usually introduces a layer of complexity to the analytical problem. In practice, however, the scale of complexity varies substantively between modelling alternatives.

The Simulated Population

7. The literature can broadly be classified into three classes concerning the in-scope population: selected individual cases; a population group (usually cross-section) at a given point in time; and a population group through time.

a. Case studies help to clarify the practical implications of alternative reform scenarios by focussing upon limited "types" of individual. Case studies, however, are ill-suited to exploring population-wide effects of tax and transfer policies, as such effects require consideration of a very broad range of alternative case-specific contexts.
b. A static microsimulation model is designed to generate individual level data for a nationally representative population at a single point in time. The micro-detail permits a very broad set of distributional analyses to be conducted, while the representative nature of the considered population permits consideration of population averages and the budgetary impact of policy alternatives.
c. A dynamic microsimulation model projects the individual units of a reference population through time, thereby allowing the analyst to consider the effects of policy into the future. Static ageing, the simplest approach to projecting the reference population through time, is based on re-weighting each micro-unit in the reference data set to reflect anticipated temporal variation. Dynamic population ageing projects each micro-unit through time, thereby generating a panel dataset (as opposed to a set of repeated cross-sectional snap-shots). This approach facilitates longer-term projections, including analysis of lifetime implications of policy alternatives, but poses serious computational challenges that have limited its use in the literature.

8. We conclude this review of existing methodologies by recommending the development of a *static structural microsimulation model* to meet the Department's objectives. A static model reduces analytical complexity, which permits greater accommodation of population heterogeneity. Coupled with a structural model of labour supply, this modelling approach is suitable for exploring the short-to-medium term effects of policy counterfactuals on labour supply for a broad subset of the population.

Roadmap to Model Delivery

9. CAPITA is currently planned to be a static microsimulation model with exogenously imposed behaviour (the default assumption being no behavioural response). This model will project the effects of selected taxes and income support payments, starting from a reference data set described by the Survey of Income and Housing (SIH).
10. We discuss how this model could be transformed into a static structural microsimulation model. We define:
    a. the inputs from CAPITA required for the behavioural module

b. the practical functioning of the behavioural module
c. the type of outputs that the module will be able to produce
d. advisable processes to adopt for validating the behavioural module
e. considerations in relation to delivery of the behavioural module

11. The inputs required by the behavioural module can be obtained either directly from the SIH, imputed from alternative data sources, or generated by the non-behavioural component of CAPITA. The report outlines the various techniques that could be employed to deliver the data that are necessary, including reference to likely data sources.
12. The behavioural module that we outline is now standard in the literature, and we provide brief details concerning its practical operation.
13. The production of individual-level output will allow the user to perform any post-simulation computation that they wish, providing maximum flexibility. For ease-of-use reasons, the module should also allow a few standard outputs to be produced at the push of a button. In this regard, the report makes specific recommendations, such as the cost and revenue implications of the policy reform or the transition matrices between the selected discrete hours levels.
14. There are three separate aspects of module validation, which are discussed at length in the report:
    a. to ensure that the programming code used to implement the module is free from error
    b. to verify that the parameters upon which the module depends have been determined/selected in a justifiable fashion and represent a sensible set of assumptions in context of the existing literature
    c. to ensure that the module is capable of generating plausible behavioural responses to policy counterfactuals
15. Delivery of the module requires attention to ensure that it is effectively used. In this regard, we make important suggestions concerning documentation, training and user support.

Two Non-Technical Case Studies

16. The report discusses how the standard model of labour supply could be adapted to allow for childcare and activity requirements. These "case studies" provide useful practical examples of the issues involved, in contexts that the Department has identified as being of particular interest.

1. Introduction

Understanding behavioural responses to the policy environment is complicated, and an entire field of economic research has been devoted to exploring the effects of policy change on employment and labour market decisions. Anyone seeking to develop analytical capacity in this field is presented with a bewildering array of possibilities, and the first challenge is to select an analytical tool that is fit for purpose. Choose a tool that is too simplistic, and subsequent analysis risks being criticised for lacking realism. Choose a tool that is too complex, and subsequent analysis risks being criticised for being unintelligible. One of the primary objectives of this report is to set out the alternatives that represent current best practice in the analysis of labour supply responses to transfer policy, including discussion of their respective advantages and disadvantages. The report also provides clear recommendations for model development based on evolving analytical capacity within the Department, and associated policy needs.

An important back-drop to this report is current modelling work being conducted jointly by the Treasury, the Department of Social Services, and the Department of Employment. These departments are now in the process of building a static microsimulation model of taxes and transfers. The new analytical tool – which is named the Commonwealth Tax and Transfer model, CAPITA – is expected to be fully operational by the end of 2014. Notably, CAPITA will hold labour supply behaviour fixed when simulating policy alternatives. This feature of the model limits its use for exploring the implications for employment of policy alternatives, which are of particular interest and relevance to the Department of Employment (hereafter referred to as "the Department").

Microsimulation models like CAPITA are widely used to explore the distributional and budgetary consequences of changes to tax and benefit regimes. This type of model is very valuable for distinguishing the effects of policy changes on households of different types. For example, at budget time we are used to statements like "A family with two young children will be better off, but a pensioner household will be worse off".

Of course, the Commonwealth has access to a suite of modelling frameworks that extend beyond the capabilities that will be available through CAPITA. The Commonwealth has, for example, access to:

- STINMOD – a non-behavioural microsimulation model similar to CAPITA, developed at NATSEM

- MITTS – a microsimulation model that includes labour supply projections based on utility maximisation theory, developed at the Melbourne Institute
- STINMOD-B – a behavioural add-on to STINMOD developed by Treasury following experience gained in the development of MITTS

This report considers the issues involved in extending the non-behavioural variant of CAPITA to permit consideration of labour supply responses to policy alternatives (with which Treasury has some previous experience). We also discuss the challenges associated with adapting such an extension to account for selected issues raised in the contemporary policy debate.

We conclude that it is feasible to extend CAPITA to permit consideration of labour supply responses to the policy environment, through the development of an add-on behavioural module (hereafter referred to as 'the module'). The report sets out sensible terms of reference, explores viable modelling alternatives in context of the contemporary literature, and considers the issues involved in practical implementation. The primary objective of this detail is to facilitate on-going development of model capacity within the Department.

The remainder of the report is organised as follows. Current best practice approaches for modelling the labour supply effects of tax and transfer policy are reviewed in Section 2, at the end of which we set out our modelling recommendations. Section 3 describes a reasonable set of conditions that should be required of the module, and Section 4 discusses two specific case studies of special interest to the Department for extending standard labour supply modelling in a microsimulation context.

2. Modelling Methodology

There is no single "best" approach for exploring labour supply responses to tax and transfer policies. A range of current best-practice methods currently exists, each with relative advantages that make them suitable for selected policy question(s). There are two key criteria over which such studies vary: the description assumed for individual behaviour, and the dimensionality of the considered population. This section summarises the alternative possibilities concerning these two criteria that have been applied in the contemporary literature, and concludes with our recommendations in view of Departmental needs.

2.1 Behavioural assumptions

The labour supply literature spans a diverse spectrum of alternative conceptions concerning individual decision making. At one extreme of this spectrum are studies that exogenously impose employment outcomes. At the other extreme are studies based upon formal models of decision making.

2.1.1 Imposed behaviour

Studies that exogenously assume labour supply typically make broad generalisations that are easy to document. Common assumptions of this type include the proposition that labour supply does not respond to policy changes, that all men work until state pension age, or that employment rates increase by a given number of percentage points in context of a counterfactual economic environment. This type of analytical approach has the advantage that it is highly transparent and readily understood. It is also relatively easy to conduct associated sensitivity analysis, by altering the exogenously assumed behaviour to reflect alternative plausible outcomes. The approach is, however, open to criticism due to the arbitrary nature of the assumed employment outcomes. It also does not provide any predictive information concerning how labour supply will respond to alternative policy contexts, as it relies entirely upon analyst views in this regard.

As noted in the introduction, the new Commonwealth microsimulation model CAPITA will assume that behaviour is invariant to the policy environment. We also return in Section 2.2.1 to discuss an analysis by Dalton and Ong (2007), who explore the implications of a change in transfer policy using exogenous assumptions concerning labour supply.

2.1.2 Reduced-form regressions

The simplest approach for allowing historical behaviour to inform views concerning the responsiveness of employment to policy alternatives is commonly referred to as reduced-form regression analysis. Reduced-form regressions are designed to summarise in an accessible form correlations between observable characteristics that are described by survey data. This methodology includes a range of alternative approaches, from simple linear equations that describe a single endogenous variable as a function of a set of hypothesised exogenous variables, through to a system of equations. The common feature that distinguishes this methodological approach is that the parameters of the equations that are considered for analysis may themselves depend upon factors that are determined by policy makers.

Consider, for example, a simple reduced-form analysis that describes labour supply decisions as a function of each individual's age and hourly wage rate. It is reasonable to suppose that any estimated coefficient on the hourly wage rate will be affected by the prevailing tax and transfer policies. This is particularly likely if gross-wages, rather than post-tax and benefit wages, are considered for analysis because high marginal tax rates drive a wedge between the impact that an hour of work has on gross incomes and on take-home pay.

Hérault *et al*. (2014) use reduced-form regressions to describe labour supply responses to changes in Australian tax and transfer policy. Using micro-data covering the period 1994 to 2010 – and hence multiple changes in policy parameters – this study identifies associations between policy parameters, and extensive and intensive margins of employment (the decision of whether or not to work, and how many hours to work). In their regression specifications, they include benefit levels as well as withdrawal rates and thresholds. One of the main results that they report is that a $1 per week increase in net income support eligibility is associated with reduced odds of employment of between 0.1% and 0.6%.

Reduced-form analysis can be a powerful tool for exploring behavioural responses to historical policy reforms. The approach can also indicate the impact of tax and transfer policy on household and government budgets. Furthermore, the approach gives the analyst a great deal of flexibility when selecting the analytical specification. This allows a large set of variables to be included in the analysis, which might be difficult to distinguish within a theoretical model of agent behaviour. This feature allows the analyst to tailor their study to the data that are available.

Nevertheless, the limitations of reduced-form analyses are widely understood by the economics profession, especially since Lucas published his critique of the approach in 1976 (Lucas, 1976). The fundamental problem is that the potential sensitivity of the parameters of a reduced-form model to the policy environment makes such models inappropriate for considering likely behavioural responses to policy counterfactuals. This has led to the development of so-called "structural models", which are explicitly designed to explore likely behavioural responses to alternative possibilities for policy reform.

2.1.3 Structural models

A structural analysis is founded on a theoretical framework that is assumed to be structurally stable, in the sense that it remains invariant to changes to the policy environment. The theoretical frameworks that are commonly adopted for analysis depend upon a set of assumptions concerning how decisions are made, which are operationalised by functional forms based on observable characteristics. Importantly, the parameters of these functional forms – referred to as "deep parameters" – are assumed to be stable in context of both temporal and policy variation. These assumptions permit the model to predict behaviour in the future, based upon behaviour observed in the past. The term "deep" here indicates that the respective parameters represent fundamental building blocks of human behaviour upon which the analysis depends.

The most common theoretical framework used to describe behaviour in economics assumes that people make decisions as if to maximise expected lifetime utility, given their preferences, constraints, and expectations about the future. Importantly, it is common to assume that expectations are rational, in the sense that they are entirely consistent with the decision-making context.

The key advantage of a structural model is that it provides a coherent basis for exploring behavioural responses to policy counterfactuals. The central draw-back of the approach is that the theoretical framework upon which the analysis depends usually introduces a layer of complexity to the analytical problem. The scale of complexity associated with a theoretical framework varies substantively between modelling alternatives. In context of the utility framework, for example, some sets of assumptions permit closed-form solutions to be derived for endogenous decisions. In this case a structural model's parameters can be estimated and used to predict behaviour in a similar way to the reduced-form models that are discussed above. However, not all models based on the utility framework have closed-form solutions, in which case numerical solution methods need to be employed. This is commonly the case, for example, in dynamic contexts subject to uncertainty.

A key criticism of the utility maximisation framework for projecting reactions to policy counterfactuals is that such frameworks over-state the flexibility of behaviour in practice. In many practical contexts, it is reasonable to suppose that there exist non-trivial rigidities that tend to dampen behavioural responses to changes in incentives. In such cases, behavioural responses tend to be muted in the short term, and the full effects of a given policy reform will only become evident after people are given some time to adapt to their altered incentives. Any model that omits such rigidities – and most models based on utility maximisation do – will consequently tend to overstate short-run behavioural responses (implying, for example, larger wage elasticities than observed in practice).

Some models based on utility maximisation do recognise the potential importance of behavioural rigidities, and build these into the decision making process (e.g., Lucchino and van de Ven, 2013). Furthermore, we note that even if interest is confined to the very short-term, models based on the utility maximisation approach that omit decision making rigidities remain useful. The reason is that the framework provides a coherent basis for identifying the implications for behaviour of changes to incentives implied by a policy counterfactual – incentives that can be very difficult to infer without use of a model. Seen in this way, someone who cares only about the very short-term might choose to discount the behavioural projections derived from a utility maximisation framework without rigidities, with more discounting applicable as the desired time horizon shortens.[1] Nevertheless, the direction of behavioural responses, and the relative size of these responses for people in different circumstances, should remain of keen interest.

Structural models are now frequently used to assess the implications of policy reform. A prominent example of this modelling approach in the Australian context is the Melbourne Institute Tax and Transfer Simulator (MITTS), which we discuss at length in Section 2.2.2.

[1] This approach has been applied in the past when using MITTS to project labour supply changes for proposed policy changes over a limited fixed number of years.

2.2 Dimensionality of the Simulated Population

The literature can broadly be classified into three classes concerning the in-scope population: selected individual cases; a population group (usually cross-section) at a given point in time; and a population group through time.

2.2.1 Individual case studies

One way to simplify analysis is to focus on a selected set of individual "cases".[2] In regards to tax and transfer policy, the case is usually defined as a person with a selected set of characteristics, and analysis focuses upon the financial implications of the policy of interest either by ex-post observation or ex-ante imputation. As such, the case study approach is capable of considering both existing policies and the implications of counterfactual policy reforms, indicating the ways in which individual circumstances interact with the policy environment.

Dalton and Ong (2007), for example, explore the effects of the Australian Welfare to Work reforms using data for existing Disability Support Pension (DSP) recipients. They consider the case of a partnered male public renter on DSP who had been economically inactive since 2002. His partner was working part-time and receiving Newstart Allowance (NSA) at a reduced rate. Dalton and Ong (2007) used a predicted wage for this individual, to consider the impact of work on their disposable income both prior to and following the Welfare to Work reforms. They find that the individual received 41 per cent higher disposable income if they worked prior to the reforms, falling to a 25 per cent premium following the reforms. The authors consequently suggest that the Welfare to Work reforms generated weak incentives to work for such people, particularly after work-related expenses are taken into account.

The analysis reported by Dalton and Ong (2007) provides a nice example of the essential mechanics that underlie the case study approach:

(1) define your subject of interest (including all individual-specific characteristics)
(2) make exogenous assumptions concerning behaviour[3]

---

[2] A "case study" is some times also referred to as a "cameo"; see, e.g. "NATSEM Budget 2014-15 Analysis", hypothetical modelling; viewed 30/09/2014 at: http://www.natsem.canberra.edu.au/storage/2014-15%20Budget%20Research%20Note.pdf

[3] It is not very common to project behaviour in step (2) listed here, despite being conceptually possible.

(3) calculate budgetary implications, given (1) and (2)

By focussing upon limited "types" of individual, case studies can help to clarify the practical implications of alternative reform scenarios. This same feature of the case study approach, however, makes it ill-suited to exploring population-wide effects of tax and transfer policies, as such effects require consideration of a very broad range of alternative case-specific contexts. This has led to the development of microsimulation models representing broad segments of the population.

2.2.2 Static microsimulation

Microsimulation models were first applied in an economic context by Orcutt (1957), and are now commonly used for analysing government policy. In these models, each micro-unit (also referred to as agent) from a population is individually represented. This facilitates analysis of heterogeneity and diversity within the simulated population. As such, microsimulation models are particularly useful for policy analyses where the effects depend upon individual-specific circumstances, or where the distributional implications are a focus of interest.

A static microsimulation model is designed to generate output relevant for a single point in time. Variants of this approach that have been adapted to consider the effects of tax and transfer policy commonly start with a reference data set that describes relevant circumstances for each benefit unit sampled to reflect a population cross-section. The information contained in the reference data set is combined with assumptions concerning agent behaviour, and the rules and regulations of an assumed transfer system to generate various measures of interest, including income gross and net of transfer payments, net tax take, rates of employment, and so on.

As discussed in Section 2.1, the simplest modelling approach is to assume that behaviour remains invariant to changes to the policy environment. In Australia, this approach was adopted for the Static Incomes Model (STINMOD) developed at NATSEM, and is the approach currently scheduled for implementation in the CAPITA model. At the other end of the behavioural spectrum, the Melbourne Institute Tax and Transfer Simulator (MITTS) developed at the Melbourne Institute is based upon a structural model of labour supply.[4]

[4] Prominent examples of static microsimulation models in the international literature include, POLIMOD (UK; see Redmond *et al.* 1998), EUROMOD (15 member states of the European Union; see Sutherland, 2001), TRIM2 (US; see Giannarelli, 1992), SPSP (Canada; refer to Statistics Canada), SWITCH (Ireland; see Callan *et*

As the theoretical model of labour supply decisions that is implemented in MITTS is what makes the model an appropriate basis for predicting labour supply responses to policy counterfactuals, we provide some further detail of it here.

The behavioural component in MITTS is based on the assumption that individuals and couples maximise utility, which is represented as a function of income, and of leisure and home production time.[5] Individuals are assumed to balance income, and hours of leisure and home production, where more of one implies less of the other. The model allows for different preferences for income and leisure/home production time, depending on the individual's characteristics. Couple families are assumed to maximise a joint utility function, and to determine their hours of work jointly. The estimated preference parameters drive behavioural responses to policy changes in MITTS.[6]

MITTS has been used extensively to analyse the effects of policy reforms. For instance, Creedy *et al.* (2009, 2011) examine the implication of abolishing the tax-free threshold on labour supply. Cai *et al.* (2008) and Duncan and Harris (2002) investigate the effect of welfare reforms on the labour supply of sole parents. More recently, Hérault *et al.* (2014) use MITTS to examine the responsiveness of labour supply to (i) a reduction in income support withdrawal rates, (ii) an increase in the withdrawal-free area and (iii) an increase in benefit payment rates. They find that lower payments tend to increase employment, particularly at the bottom end of the earnings distribution and to a lesser extent for moderate and high wage earners. Payment rates are shown to be qualitatively more important for employment outcomes in aggregate than withdrawal thresholds and withdrawal rates.

Static microsimulation models consequently provide a potentially powerful tool for exploring how policy change is likely to affect a population at a given point in time. By treating each

al. 1996), LOTTE (Norway; see Fjærli *et al.* 1995), and FASIT (Sweden; refer to the Swedish Ministry of Finance).

[5] Total time available is assumed to be divided between leisure and home production time on the one hand, and hours in employment on the other hand. Leisure and home production time are usually combined given the data that are available.

[6] Creedy *et al.* (2002) outline the initial set-up of MITTS. Creedy and Kalb (2006) discuss the methodology of behavioural microsimulation modelling more generally, including a number of examples using MITTS, whereas Buddelmeyer *et al.* (2007) report on a range of microsimulation applications in tax policy design.

micro-unit of a reference database as a separate case study, microsimulation can be used to build up a macro-level picture of the effects of a change to transfer policy. The micro-detail permits a very broad set of analyses to be conducted, from the consideration of population averages and budgetary impact of policy alternatives, through to distributional effects. Static microsimulation models do not, however, allow the analyst to consider the effects of policy through time – for that we need a dynamic microsimulation framework.

2.2.3 Dynamic microsimulation

Dynamic microsimulation models are designed to project a reference population through time. Starting from a static microsimulation model, the simplest approach to projecting the reference population through time is to re-weight each micro-unit in the reference data set to reflect anticipated temporal variation. This approach is commonly referred to as "static ageing", and is usually considered appropriate for projecting population circumstances over limited time-horizons (5 to 10 years).

The problem with using static ageing to project over longer time frames is that it becomes increasingly difficult to infer how a change in tax and transfer policy will influence the relative weighting of alternative micro-units reported in the reference data set. Suppose, for example, that a change in tax and benefits policy had the effect of increasing labour supply. Increased labour supply over a short time horizon can be expected to have a limited impact on the underlying circumstances of individuals including the wages that they are eligible for and the savings that they have. Over longer time horizons, however, both wages and wealth are likely to be influenced by a sustained increase in labour supply, complicating attempts to specify the weighting adjustments required for static ageing.

Furthermore, static ageing does not permit analysis of the effects of a policy reform on individuals through time. A model based on static ageing therefore cannot provide information about how a policy is likely to influence alternative individuals during the course of their respective lives.

Dynamic population ageing is designed to address the two key problems identified above for static ageing. Microsimulation models based on dynamic ageing track each individual described by the reference database through time, thereby building up a panel dataset. For example, most dynamic microsimulation models based on dynamic population ageing that are designed to consider the effects of transfer policies start with a reference database that describes marital status, parenthood, employment status, income, and previously accrued measures of wealth. The population is then aged one year with reference to a series of transition equations that describe the dynamic evolution of individual-specific circumstances. Where these equations describe stochastic events, random draws are taken from the underlying statistical distributions in a process that is commonly referred to as Monte Carlo simulation. The income of each individual in each time period, for example, is usually simulated based on characteristics such as the individual's past income, their accumulated experience, demographic characteristics, and on a random draw from a log-normal distribution that accounts for unexplained variation. By repeating this procedure over successive years, the life history of each individual in the reference database can be generated.

Dynamic population ageing introduces two important complications, relative to either static microsimulation models or statically aged models. The first is the need to describe how individual-specific circumstances evolve through time. In most practical contexts, the evolution of individual-specific circumstances depends on individual decisions, thereby exaggerating the importance of a structural analysis. The second important complication is that adopting a structural approach to behavioural simulation is generally hampered in a dynamic context by the non-existence of closed-form solutions to the decision problem (discussed in Section 2.1). A series of important papers[7] have shown that short-cuts designed to address the challenges posed by non-existence of closed-form solutions can result in substantial distortions, suggesting that computationally demanding numerical solution methods need to be employed. The trade-off when adopting such methods is that, in context of contemporary computing technology, such models cannot account for the diversity of individual-specific circumstances that is present in most population cross-sections (and which are commonly used in static microsimulations).

The complications associated with dynamic ageing have limited its use in the contemporary literature. Most existing models of this type have mitigated the difficulties involved by adopting reduced-form approaches to modelling behaviour, or stylised assumptions that permit closed-form solutions to the decision problem. This is in spite of the well-recognised risks associated with such approaches (discussed above). Examples in this vein for Australia include

[7] See especially Kimball (1990), Deaton (1991) and Carrol (1992); a simple worked example is provided by Browning and Lusardi (1996).

the reduced form (non-behavioural) model for selected birth cohorts considered by Creedy and van de Ven (1999), HARDING (Harding, 1993), and DYNAMOD-2 (King *et al.*, 1999). For examples from the international literature, see PENSIM2 (UK; see Curry, 1996), LIFEMOD (UK; see Falkingham and Lessof, 1992), ASPEN (US; see Basu *et al.*, 1998), CORSIM (US; see Caldwell, 1997), DYNACAN (Canada; refer to Statistics Canada, based on DYNASIM, see Orcutt *et al.*, 1976), MICROHUS (Sweden; see Andersson *et al.*, 1992), NEDYMAS (Netherlands; see Nelissen, 1994), SESIM (Sweden; refer to the Swedish Ministry of Finance), Destinie (France; see Dueé and Rebillard, 2004), and SAGE (Britain; see Zaidi, 2007). Nevertheless, advances in computing power are now starting to permit development of policy-relevant structural dynamic microsimulation models that treat uncertainty as a centrally important feature underlying simulated behaviour; e.g. LINDA (UK; see Lucchino and van de Ven, 2013).

2.3 Modelling recommendations

It is our view that the objective to extend analytical capacity within the Department to permit consideration of employment responses to transfer policy would be best met by a *static structural microsimulation model*.

A *static model* reduces analytical complexity, thereby accommodating the full heterogeneity of the population under consideration. If there is a need to consider the influence of transfer policy on employment outcomes into the future, then static ageing can be used to project the reference population through time (albeit over a period that does not exceed 10 years).

A *structural model* is required to permit the model to explore the effects of policy counterfactuals on labour supply. The utility framework that represents current best practice in structural analysis of private individual decision making is well established in the existing economic literature, which limits the risks involved in associated model development. This modelling approach can be extended to explore labour supply decisions by a broad subset of the population, including those with health issues or close to retirement age.

A *microsimulation model* is sufficiently flexible to explore population-wide effects, rather than the implications for isolated population subgroups or hypothetical individuals identified by a limited number of characteristics that are usually the focus of analysis of alternative analytical approaches.

This conclusion is further informed by the joint project by the Treasury, the Department of Social Services, and the Department of Employment which is currently underway to develop CAPITA, a static microsimulation model. It is possible to extend this model, which is currently scheduled to be completed by the end of 2014, to consider policy implications for employment through the development of an add-on module. By taking advantage of existing analytical capacity in this way, the approach promises to reduce substantively the costs that are commonly associated with development of a static structural microsimulation model that is suitable for exploring employment responses to alternative policy interventions. The steps involved in developing such a module are described in Section 3.

3. Roadmap to Module Delivery

3.1 Module Development

This section discusses the inputs from CAPITA required for the behavioural module, the practical functioning of the module, and the type of output that the module will be able to produce. Adding a behavioural module to CAPITA requires the establishment of a communication process between the module and CAPITA. Although in principle this communication process can be either two-way or one-way, we have framed the following discussion in terms of a one-way process for clarity and without loss of generality.

3.1.1 Inputs

A number of variables required for the behavioural module can be extracted directly from the Survey of Income and Housing (SIH), the base data source used for CAPITA. Other variables have to be generated by CAPITA.

Figure 1 shows the input required for the behavioural module. Individual and tax unit demographic characteristics – including tax unit detail such as size, composition and location, and individual level detail such as age, health and education status – are needed so that the module can allow behavioural variation based on differences in each of these characteristics. These data could be obtained directly from the SIH. It would be useful, however, for CAPITA to export these data to the module, to ensure that the module is fully consistent with the wider model framework.

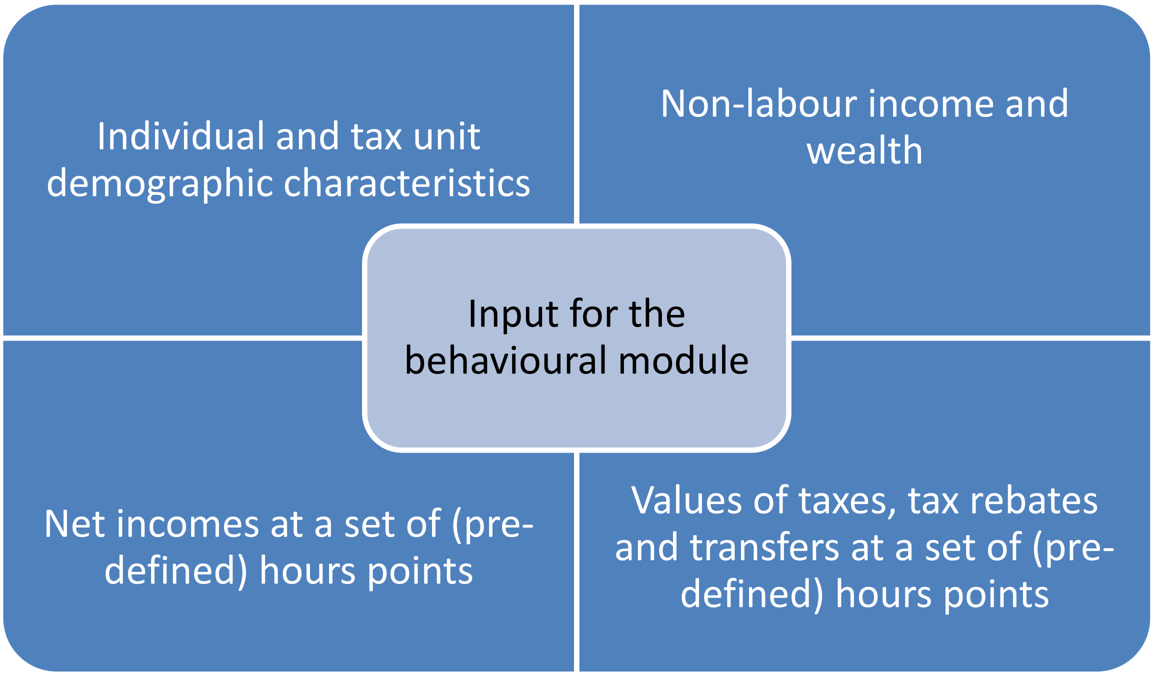

Figure 1 Required input for the behavioural module

Information on health is required to determine eligibility to some benefit payments, as well as distinguish how labour supply behaviour varies with health status. Unfortunately, the SIH contains very limited information concerning health status. If health related benefits, and the bearing of health on labour supply, are to be included in the module, then each individual's health status must be imputed from secondary data sources.

One approach that could be used to impute health data into the module is to estimate a statistical model for health on data derived from the Australian Health Survey (AHS).[8] This statistical model for health status would reference a set of individual-specific characteristics that are reported in both the SIH and the AHS. A multi-level ordered logit model could, for example,

[8] This will be possible if Confidentialised Unit Record Files (CURFs) for the AHS are made available by the Australian Bureau of Statistics (ABS). The ABS website only reports Tablebuilder under the available microdata related to the 2011-12 AHS. However, under the expected microdata, two AHS-related files are listed as expanded CURFs which will become available in the second half of 2014. They are *Australian Health Survey, National Health Survey with Biomeds, 2011-12*; and *Australian Health Survey, Core Content - Risk Factors and Selected Health Conditions, 2011-12*. Hopefully, these will contain what is needed to estimate health status through the Remote Access Data Laboratory. Alternatively, the Survey of Disability, Ageing and Carers (SDAC) or the Household, Income and Labour Dynamics in Australia (HILDA) could be used.

be used to describe the statistical propensity of alternative health conditions – distinguished by type (mental/physical) and severity – based on selected individual characteristics, such as age, ethnicity, relationship status, labour force status, income and so on. Given these propensities for each individual, a random number draw could be used to assign health status for inclusion in the computations associated with both eligibility for income support and employment responses.

Income information is crucially important to the module. The behavioural module requires CAPITA to compute pre-tax and transfer incomes as well as incomes net of tax and transfer payments for a set of pre-determined hours points. The module will organise these data into a vector of net incomes for each tax unit, together with the corresponding values of taxes paid, tax rebates, and transfers received. For individuals of working age, but not observed to be working, an hourly wage rate would be imputed using a wage equation that is internal to the module structure. Hence, for each tax unit, we can obtain the precise budget constraint establishing the relationship between hours worked and net income. These are the basic data that the behavioural module uses to evaluate which part of the individual's budget constraint is optimal (see section 3.1.2 for details).

Any attempt to address issues around retirement decisions requires information on wealth and non-labour income. Fortunately, wealth data have been collected as part of the SIH in three years, the most recent being the 2009/10 wave of the survey. It is recommended that the behavioural module be used strictly in conjunction with SIH data that include wealth data whenever simulation of labour supply late in working life is a subject of particular concern.

3.1.2 The functioning of the behavioural module

The behavioural module will take the inputs described above from CAPITA and produce the relevant labour supply responses, which give rise to a set of outputs as described in section 3.1.3. In practice, the module will be based on a set of pre-estimated preference parameters. In other words, it will be a quantifiable behavioural model of individual tastes for net income and labour supply, which can be used to simulate individuals' preferred labour supply under alternative sets of economic circumstances.

Typically, discrete choice structural labour supply models are chosen for their ability to accommodate complex and non-linear budget constraints (i.e., relationships between hours worked and net income). The use of direct utility functions is also generally preferred as it means that no integration from estimated supply functions is required, which simplifies the simulation process considerably. In addition, modelling (and predicting) transitions into and out of the labour market is relatively straightforward in the discrete model compared to the continuous model.

Utility is maximised conditional on constrained resources: that is, the total amount of time available to each adult and the total amount of household income an individual is able to earn through labour force participation. It is expected that utility increases with an increase in leisure (including home production time) and income. Usually less leisure time means more income, except in context of marginal tax rates equal to or greater than 1. Explicit assumptions are required concerning take-up rates for the benefits that an individual is eligible for. Maximising a household's utility involves balancing the amount of leisure and income. Typically, it is assumed that all non-participants in the labour market are voluntarily not working and that participants are at their preferred labour supply points. Wage rates, non-labour income and household composition are exogenous.

Utility is assumed to consist of a deterministic and a random component, reflecting unobserved heterogeneity. Choosing an extreme value specification for the random component results in a multinomial logit model; see Maddala (1983) and Creedy and Kalb (2005). Depending on the form of the utility function, different interactions between household income and labour supply can be modelled. A quadratic specification for the deterministic component is standard in the associated literature (e.g. Keane and Moffitt, 1998). It is simple but flexible in that it allows for leisure and income to be (direct) substitutes or complements. A fixed cost of working parameter is included with the income variable to indicate the cost of working versus nonparticipation, following Callan and Van Soest (1996).

Simulations are probabilistic, as utility at each discrete hours level is specified as the sum of a deterministic component (depending on the hours worked and net income) and a random component. Hence the behavioural module does not identify a particular level of hours worked for each individual after the policy change, but generates a probability distribution over the discrete hours levels used. Given a random set of draws from the error term distribution, along with the computed deterministic component of utility at each of the labour supply points, the optimal choice for each draw can be determined conditional on the relevant set of error terms.

The behavioural simulation begins by recording the discrete hours level for each individual that is closest to the observed hours level. Then, given the parameter estimates of the preference function, which vary according to a range of characteristics, a set of random draws is taken from the distribution of the error term for each hours level for those whose labour supply is allowed to change. The distribution of the random draws can be made conditional on the discretised optimal hours level observed before the reform (see Bourguignon *et al.*, 1998). That is, the error terms drawn from this distribution all place the individual at the observed labour supply before the reform. This process is described as 'calibration'. A user-specified total number of draws is produced.

For the post-reform analysis, the new net incomes cause the deterministic component of utility at each hours level to change, so using the set of draws from the calibration stage, a new set of optimal hours of work can be produced; one for each draw. Thus the same random utility components are used before and after the reform. This gives rise to a probability distribution over the set of discrete hours for each individual under the new tax and transfer structure. Transition matrices showing probabilities of movement between hours levels (and movements in and out of the labour market) are computed using these probabilities.

It is straightforward to incorporate the computation of confidence intervals for the behavioural responses. However, this computation is very time consuming, so it is recommended to make the inclusion of this computation optional. Users can decide whether they need confidence intervals, and experiment with alternative simulations before doing a final version that includes calculation of confidence intervals.

The behavioural module set out above is capable of doing more than projecting labour supply. With appropriate specification, it could, for example, be used to account for the effects of activity tests and mutual obligations. This would require activity tests to be translated into welfare costs, a task that could be achieved in a similar fashion to the way that any other (unobservable) preference parameter is identified in the existing literature.

For example, suppose that the Government required everyone on Newstart Allowance to work for at least 20 hours per week on a community project. One way to allow for this requirement in the module is to alter the decision alternatives considered for analysis. In the absence of the mutual obligation referred to above, we might assume that any individual who does not work receives income support. Including the mutual obligation might introduce two alternatives: to supply no effort and receive no income support, or to supply 20 hours of effort and receive income support. Depending on an individual's market wage, outside options may become more attractive relative to income support in this context. In addition, working on the community project may be less attractive than the type of job they could obtain in the labour market given their qualifications. This would create a disutility from the 20 hours of community work that is larger than 20 hours worked in an alternative free market job.

To give another example, suppose the government introduced additional requirements for individuals in receipt of Newstart Allowance to attend the Centrelink office. This would increase the personal cost attached to receipt of income support ('cost' interpreted broadly here to account for intangible factors such as social stigma), which could be represented in the structural model by an increase in the disutility associated with receiving income support. Disutility derived from receipt of income support in the base case (our observed data) can be estimated as part of the labour supply model, where individuals choose a number of hours to work and whether or not to take up income support when they are eligible (e.g. see Moffitt, 1983; Hoynes, 1996, Keane and Moffitt, 1998; and Kalb, 2000 for Australia). Once estimated, the level of disutility associated with receipt of income support could be exogenously varied to reflect alternative policy counterfactuals.

Note, however, that the qualitative nature of mutual obligations frustrates any attempt at delivering precise measures of the disutility attached to individual policy initiatives. Exploring policy alternatives concerning mutual obligations will require arbitrary assignments of disutility, suggesting the need for extensive sensitivity analysis. Such exploratory exercises should therefore be considered as sophisticated thought experiments, rather than formal behavioural predictions. We return to discuss related issues at greater length in Section 4.2.

3.1.3 Outputs

The behavioural labour supply module to be attached to CAPITA should be required to generate an output data set containing all of the simulated detail for each individual and household in the base sample. This fundamental output will allow the user to perform any post-simulation computation that they wish, providing maximum flexibility. Furthermore, for ease-of-use reasons, the module should also allow a few standard outputs to be produced at the push of a button. We make recommendations concerning these "standard outputs" here:[9]

[9] All results can be weighted using the population weights, so predicted results can easily be obtained at the population level.

1. At the end of each simulation, summary information regarding the cost and revenue implications of the tax reform will be available, along with aggregate information about hours worked and changes in hours worked. This information will be available in tables that are automatically produced and can be exported for use in reports and other documents. If the user has selected the option to compute confidence intervals before running the behavioural simulation, then these will be provided in the summary table as well.[10]
2. Upon request transition matrices, which show probabilities of movement between the selected discrete hours levels before and after the reform, can be produced separately for each of the demographic groups (e.g. single men, single women, single parents, partnered men, and partnered women) in turn.
3. Upon request the labour force participation effects of the simulated tax reform can be tabulated for a number of categorical variables.[11] In addition, characteristics of example individuals falling into a specific category of participation effect can be retrieved for closer examination.
4. Similarly, distributions of changes in hours can be obtained for the categories listed in footnote 9. And again, it should be possible to examine specific individuals from the basic dataset falling into a specific category of change in hours.
5. Implications of the reforms for the distribution of income, taking into account the changes in labour supply, should also be made available upon request. In this regard, we recommend the method described in Creedy, Kalb and Scutella (2004), which allows for the probabilistic nature of the output. Both inequality and poverty measures can be calculated.
    a. The available inequality measures should include: Lorenz curve and Gini coefficient, Atkinson index, Theil's entropy measure, Coefficient of variation, 90/10 ratio (which refers to the ratio of income levels at the 90th and 10th percentiles of the population-weighted net income distribution), and measures of the distribution of employment.

[10] Note that the computation of confidence intervals tends to lengthen the time required to run the simulation considerably.

[11] For example, these variables could include: 1 Household income decile, 2 Income unit type, 3 Household type, 4 Family type, 5 State of residence, 6 No. of dependent children in unit, 7 Age of youngest child, 8 Gender, 9 Tenure type, 10 Marital status, 11 Country of birth, 12 Highest educational qualification, 13 Study status, 14 Age, 15 Employment status, 16 Occupation in major job, 17 Industry of major job, 18 Dwelling type, 19 Health status, 20 Long-term health condition or disability, 21 Wealth.

b. The available poverty measures should include: the Headcount which is the proportion of the population found on or below the poverty line; the TIP curve which refers to three characteristics of poverty (the Three "I"s of Poverty.) - its incidence, intensity and inequality[12]; the Foster, Greer and Thorbecke poverty index[13], and the Sen index, which incorporates a measure of inequality among the poor.
c. The user should have some choice in the inequality and poverty settings: 1) an equivalence scale can be selected (e.g. OECD or Whiteford scale), 2) the unit of analysis can be selected (e.g. income unit, adult equivalents, person), and 3) before selecting poverty measures, the poverty line can be specified (e.g. 50 per cent of median income, 50 per cent of mean income, or alternatively, instead of 50 per cent, 40 or 60 per cent could be chosen for example).[14]

6. Implications of the simulated policy changes for social welfare, again taking into account labour supply changes, following the approach proposed by Creedy and Kalb (2005).[15] These measures take into account the value of home production/ leisure time.

3.2 Module Validation

Validation is crucially important to ensuring that the module is suitable for policy analysis. There are three separate aspects of module validation. The first is to ensure that the programming code used to implement the module is free from error. The second is to verify that the parameters upon which the module depends have been determined/selected in a

[12] It indicates not only whether an individual is below the poverty line, but also the absolute poverty gap, or the extent to which income falls below the poverty line. The TIP curve is obtained by plotting the total poverty gap per capita against the corresponding proportion of people. The slope of the non-horizontal section of the curve at any point is the poverty gap at that point. The extent to which the poverty gap falls as income rises is therefore reflected in the flattening of the TIP curve, so the extent of inequality among the poor is shown by the concavity of the TIP curve.

[13] Depending on the parameter value chosen it can equal: i) the headcount, ii) a measure which depends on both the proportion of people in poverty and the extent to which, on average, those people fall below the poverty line, or iii) a measure which additionally depends on the inequality of those in poverty, as measured by the coefficient of variation.

[14] The choice of poverty lines presented here all refer to relative poverty lines.

[15] An example of an empirical application is available in Creedy, Hérault and Kalb (2009).

justifiable fashion and represent a sensible set of assumptions in context of the existing literature. The third is to ensure that the module is capable of generating plausible behavioural responses to policy counterfactuals. Each of these three issues is discussed in turn below.

3.2.1 Validating programming code

The difficulty of ensuring the validity of programming code is governed by the complexity of the analytical problem. For a module of the type considered here, it will be challenging for any one person to guarantee that the programming code is free from error. This limitation is best addressed by subjecting the code to extensive testing, a task that will continue by default for as long as the module is used.

The problem of confirming the accuracy of the programming code can also be mitigated by ensuring that multiple people have access to, and understand the workings of the module's code. Each of these people should examine the code independently prior to project sign-off, in an effort to track down computational errors. It is advisable that at least one of these people be internal to the government departments that would be using the module. This will serve to improve the transparency of the module, in addition to mitigating associated personnel risk.

3.2.2 Validating module parameters

It is important that the parameters upon which the module depends be accessible, clearly defined, and determined/selected in a justifiable fashion. The last two of these objectives are best met by requiring that a detailed technical report to be produced, as discussed in Section 3.3 below. Concerning the first objective, ensuring that module parameters are readily accessible will aid the process of checking the module programming code. It will also facilitate associated sensitivity analysis of simulated results to parameter alternatives after the module is delivered.

3.2.3 Validating module outputs

There are two different ways in which module outputs should be checked.

First, it is important to verify that the module can successfully approximate a series of observable margins under a given "base" simulation context. A discrete set of simulation targets should be defined and agreed for the base simulation context prior to the project commencing. In the process of executing the project, it should be possible to identify whether the objective targets have been achieved. If the targets cannot be achieved, then clear and sensible reasons for the failure should be provided, and a decision made concerning whether the module is fit for use despite the failure.

Secondly, it is important to verify that the module generates sensible labour supply responses to relatively simple "counterfactual" simulation contexts. A limited number of such simulation counterfactuals should be defined and agreed prior to the project commencing. Labour supply responses to these counterfactuals should be generated and interpreted to verify that the module is working as intended.

3.3 Module Delivery

3.3.1 Documentation

One of the key risks of any microsimulation model is that the calculations underlying associated output are so opaque that the model is regarded as a black-box. This is a problem because an understanding of the relationship between model inputs and outputs is crucially important to justifying projected behavioural responses; model outputs that are hard to rationalise are usually difficult to defend. Detailed documentation is therefore vitally important to ensuring that the module provides a suitable basis for policy design and analysis.

It seems reasonable to require that a technical paper and a user guide be generated as part of the module's development. The technical paper should provide a detailed description of the theoretical framework (if relevant) adopted for analysis, the way that the module interacts with the broader model structure, and how parameters of the module have been identified. Discussion concerning module parameterisation should detail how parameters were determined, what data sources (if any) were used, associated estimation methods and accompanying statistical measures of fit.

The information provided in the technical paper may not be necessary for day-to-day use of the simulation module, but is likely to be crucial when interpreting and justifying associated output. The expectation is that an undergraduate degree in economics (with a good understanding of statistics) is likely to be necessary to read the technical paper, as some technical terminology will be unavoidable in explaining the economic modelling.

The user manual, in contrast, should describe how the module can be used in language that is accessible to a general reader. This document should provide a brief non-technical description of how the module works, including some information concerning how the module inter-acts with the broader model structure. The user manual should also describe what outputs are generated by the module, outline best-practice methods of use, and provide directions for where more detailed information can be obtained.

The documentation that is outlined above could be supplemented according to future needs, via stand-alone research agreements.

3.3.2 Training

Some training is likely to be necessary after the module has been signed off and delivered. Here we briefly describe three courses that may be of interest:

1. Structure and Operations: this is an introductory course explaining the different calculations that can be carried out using the module. The course would consist of one day in a computer lab, where the different facilities will briefly be described and participants shown how to operate them. In addition to discussion of how to operate the module, the output files generated by the module, which can be explored using standard statistical packages, will be described. Explanation of features will be alternated with small practical examples and tasks will be given to participants to provide learning by experience. At the end of the course, participants should be able to run behavioural simulations using the module and be able to interpret the results from simulations.

   The user manual described above will be an important document in this course. In addition to this, a brief overview of the different components of the module will be provided together with examples of simulations and output that can be generated by the module.

2. Econometric Estimations: this is a more technically orientated course, which explains the key subcomponents underlying the module. The training session would be divided into two broad halves, one focussing on the wage models used to impute wages for non-workers, and another focussing on the labour supply models that are used to generate labour supply responses. The main focus of the course would be to explain how key components of the module are parameterised, with some attention paid to alternative approaches applied in the literature. Knowledge of basic statistics/econometrics will be assumed, but the training could start with a brief (refresher) course on probit and multinomial logit models, which are two important empirical approaches used in structural labour supply modelling. At the end of the course, participants should understand how the module has been parameterised, including an appreciation of the assumptions that are involved.

   The material used in this course would be a selection of published papers/articles. In addition, specific training notes could be prepared that summarise the content of the course. A small illustrative exercise, estimating labour supply making use of the programmes developed by the provider, could be included as part of the course.

3. Training course on the code relating to the behavioural simulation: this would be a detailed technical course focussing upon the workings of the module's programming code. The first component of the course could be quite general, providing an overview of the different steps in the code, whereas the second component would provide a more in-depth overview of the code, focussing on the structure of the module with regard to reading files and calling procedures, and explaining how to find particular subcomponents of the module in the code. The aim of the course would be to give an overview of the different files and procedures of the module, giving special attention to the code that drives the simulated behavioural responses. The course could consist of formal training, prepared beforehand by the provider, and informal interaction with course participants, providing an opportunity for questions or discussion of specific issues in the code which may have arisen while staff have examined/experimented with the code. The course should be given in a computer lab, so that course participants can browse through the code that is discussed on their own screens. Such a course would be most useful after Departmental users have had an opportunity to work with the module themselves.

   At the end of this course, participants should have a good understanding of the different files used to implement the module, and know how to learn more about specific components of the programming code. They should feel confident about making changes to the code, so that they will be able to adapt the module's output or alter preference parameters. They should also have an understanding of the specific calculations that are carried out by the behavioural module, including (for example) the computation of confidence intervals.

   The training material would consist of prepared notes written for the course (and annotated code). Prepared examples will be central to this training component, because it is likely to be infeasible to go through the entire programming code within a few hours. The examples will discuss a general approach to making changes, including tips for debugging, which can be followed in a variety of situations.

3.3.3 User Support

The training outlined in Section 3.3.2 could be complemented by on-going user support, documentation services, and code walk-through assistance as required on demand. We discuss some of the options here.

Staff from the Commonwealth Departments using the module should have access to support staff from the provider of the module. They can either e-mail questions or call with inquiries regarding the use of the module. A logbook with the questions, the date the question was asked and the time needed to address the question will be kept. Based on the logbook, a monthly report and invoice can be generated.

In addition, it is anticipated that the preference parameters feeding into the module would need to be updated every 4 to 5 years (in line with the issuing of wealth data with the SIH) to ensure that predicted responses are informed by up-to-date information. However, given the usual stability of average preferences, it is anticipated that an update of the module could be undertaken substantially more quickly than the time taken to parameterise the module in the first instance.

4. Two Non-Technical Case Studies

The section discusses two potential adaptations of the "simple" model outlined in Section 3, which have been identified as being of particular interest by the Department.

4.1 Modelling Childcare Use

Here we discuss two alternative ways in which childcare could be accommodated within the labour supply module. First, a minimal modification is discussed in which the module is extended to take childcare costs into consideration. The second alternative adapts the module to permit consideration of behavioural responses to changes in incentives for the use of childcare.

4.1.1 A simple extension of the current behavioural model

Assume that childcare does not generate costs other than the direct financial costs levied by childcare providers. Also, assume that childcare does not generate benefits, other than relaxing constraints on hours available for parental labour supply. Finally, assume that there are no restrictions on childcare availability (as no associated public information is available). Given these assumptions, a simple two-step approach can be adopted, in which childcare costs are imputed in the first step and fed into the labour supply model in the second step.

Doiron and Kalb (2005) describe one way in which the above approach could be implemented. In their description, the first step of modelling childcare costs is implemented via a reduced-form bivariate Tobit model to account for formal care use (in hours per week) and informal care cost (in dollars per week). The costs of formal childcare can then be obtained by multiplying hours of use by average childcare fees relevant for a given household, based on the age of their child(ren) and their region of residency.

Explanatory variables for the Tobit model described above include hours of work, number and age of children, presence of other adults in the household, and price of formal childcare (from an external data source to avoid endogeneity issues due to price being a choice by the parents as well). The choice of care use is continuous, depending on the amount of labour supply. No fixed link is imposed on the relationship between hours of work, and hours of formal care and cost of informal care. Instead, the link is estimated based on observed labour supply and childcare use behaviour, allowing for unobserved heterogeneity.

A contemporary estimation of the Tobit model described above would benefit from the improved data that are now available. In recent years, the SIH has reported data for childcare hours and costs, distinguishing between formal and informal care arrangements. An external source of data to obtain average hourly childcare fees by age of the child and State of residency is the 2013 National Early Childhood Education and Care Workforce Census (Department of Education, 2014), supplemented by information from the Child Care Management System if required (and available). In calculating average fees, the hourly fees of different types of childcare are weighted by the proportion of children of a particular age using that type of childcare. For an example of how this approach would be implemented, see Kalb and Lee (2008).

In the second stage of estimation, the utility of couple families, $U(C, L_f, L_m) = \nu(C, L_f, L_m) + \varepsilon$, is maximised subject to a time constraint for each adult:

$$L_f + H_f = T \quad \text{and} \quad L_m + H_m = T \qquad (1)$$

where $\nu(.)$ is the deterministic part of the utility function and ε is a stochastic error term, C is consumption which is assumed to be equal to disposable income, T denotes total hours (per week), L refers to non-work hours and is assumed to increase utility, H refers to work hours, and the subscripts f and m denote father and mother respectively. Importantly, note that childcare use does not enter into the utility function. Childcare use is exogenous and is assumed to only affect the decision problem through the impact that it has on disposable incomes at different labour supply points, as discussed below.

Households are assumed to maximise utility subject to a standard budget constraint:

$$\begin{aligned} C = {} & w_f H_f + w_m H_m + y_f + y_m + B(hc, cc, w_f H_f + w_m H_m + y_f + y_m) - \\ & \tau(B, cc, w_f H_f + w_m H_m + y_f + y_m, hc) - cc(H_f, H_m, hc, w_f H_f + w_m H_m + y_f + y_m) \end{aligned} \qquad (2)$$

where w indicates gross wage rates (either observed or imputed); y non-labour incomes; $B(.)$ is the amount of benefit a household is eligible for given their household characteristics hc, gross childcare costs cc, and household income; and τ is the tax function that indicates the amount of tax to be paid. A similar utility function and budget constraint can be constructed for single parent families, allowing for one labour supply choice only.

The budget constraint defined by equation (2) has been adapted to accommodate childcare costs by the associated benefit $B(.)$ and tax $\tau(.)$ functions. These are elements of the static microsimulation model, e.g. CAPITA. The inclusion of childcare subsidies and rebates is complicated, however, by the alternative approaches that exist for minimising the net effective childcare costs incurred. One individual, for example, might use a salary sacrifice scheme, while another takes advantage of the Childcare Benefit and Childcare Rebate. In view of the primary objective of keeping the resulting model as simple as possible, we recommend that the calculations necessary to evaluate the impact of childcare costs on household budgets should be managed within a separate program module so it can be easily switched on and off again.

Having implemented code to project gross childcare costs, and to calculate the associated impact on household disposable incomes, simulation of labour supply proceeds in the same way as it does without childcare costs. First, for each hours level allowed in the discrete choice labour supply model, a gross income level (together with all transfers and taxes) is computed. Then, for each household with children aged 12 years or younger in the SIH, a predicted net cost of childcare is imputed at each labour supply point, based on the characteristics of the household (State, urban residency, number and age of children, and calculated gross income). This imputed value is based on the observed usage in the data, and is not derived from utility maximisation under different policy settings, so childcare use does not respond to policy changes.

A key point to note is that, independent of how the childcare cost *cc* is imputed, it is best to recognise the uncertainty of the associated projection. In context of the Tobit model described above, this can be done by drawing repeatedly from the joint distribution of formal and informal childcare costs. A simulated maximum likelihood approach can then be used to estimate the model's parameters conditional on these draws. The draws are generated by including a draw from the error term (individual-specific variation) when predicting childcare costs and demand using the model. An advantage is that the calculation of the Child Care Benefit is more accurate in this approach compared to using expected childcare costs, given that the subsidy payable for the average childcare cost is not the same as the average Child Care Benefit based on the potential different outcomes for childcare costs.

In summary, the approach to simulating the influence of childcare costs that is described in this section involves three key steps:

1) Estimate the incidence of childcare costs as a function of household characteristics including the number and age of children, and labour supply decisions
2) Calculate the gross childcare costs implied by step 1 for labour supply alternatives within the model structure.
3) Augment programming modules that calculate taxes and benefits to compute disposable income allowing for gross childcare costs identified via step 2, and for corresponding childcare subsidies.

4.1.2 Modelling childcare use decisions explicitly, and jointly with labour supply

In the modelling approach described above, childcare use is exogenously defined, there is no explicit consideration of childcare quality, or allowance for childcare to affect utility directly. This section considers how a standard labour supply model could be adapted to allow childcare use to be endogenously determined within the model structure. The central advantage of this is that the resulting model would predict the impact of policy alternatives on the demand for childcare, whereas in the approach set out above, childcare use is specified to reflect historical data.

The simultaneous choice of childcare and labour supply is a multidimensional issue, which introduces substantial analytical challenges. The associated literature has simplified this problem by focussing exclusively on childcare options for mothers, taking fathers' labour supply as given. We suggest an adaptation of this approach here.

The mother is assumed to choose the job and childcare arrangement that maximises preferences subject to a standard budget constraint. The modelling framework in this approach is based on the notion that parents choose a level of labour supply for the mother and a childcare alternative for a certain number of hours for each child. Both choices affect the household's utility level. For practical reasons (such as lack of information on availability of informal care and the diversity of childcare belonging to this category), it is assumed that informal childcare does not contribute directly to the household's utility. Thus, households maximise utility represented by $U\left(C, H_m, H_{ccf}\right)$, the utility of working H_m hours by the mother and using H_{ccf} hours of formal childcare, where C is consumption/disposable income. Preferences can be specified as:

$$U\left(C, H_m, H_{ccf}\right) = v\left(C, H_m, H_{ccf}\right) + \varepsilon, \tag{3}$$

where $v\left(C, H_m, H_{cc}\right)$ is the deterministic part of the utility function and ε is a stochastic error term. We assume the same time constraint as in section 4.1 for the mother, where L_m (implicitly) includes maternal care for the child(ren):

$$L_m + H_m = T \tag{4}$$

and, for each child j under 12, subject to a child time constraint:

$$H_{jccf} + L_{jmcc} + H_{jcci} = T - D_{jc3\text{-}4}(H_{PS}) - D_{jc5\text{-}12}(H_S) \tag{5}$$

where H_{jccf} are the hours of formal childcare for child j, H_{jcci} are the hours of informal childcare (or father care) for child j, L_{jmcc} are the hours of care provided by the mother, $D_{jc3\text{-}4}$ and $D_{jc5\text{-}12}$ are indicators for the age of the child which determine whether the child can attend preschool and school respectively, H_{PS} and H_S are the hours of preschool and school generally provided. We define a standard household budget constraint:

$$\begin{aligned} C = & \, w_m H_m + y_m + y_f + B(hc, cc, w_m H_m + y_m + y_f) - \\ & \tau(B, cc, w_m H_m + y_m + y_f, hc) - cc(\sum_j H_{jccf}\, p_{jccf}, hc, w_m H_m + y_m + y_f) \end{aligned} \tag{6}$$

where y_f now includes all income of the father (labour and non-labour), and p_{jccf} is the hourly fee of formal childcare for child j which depends on location and age of the child.

Hours of care by the mother are not observed in the data; only the total number of hours of care (and total cost) within the income unit is observed for all children together, but hours and cost are separately provided for formal and informal care.[16] Since no information on the availability of informal care is observed (when informal care is not used by the household), it cannot be treated as a choice. In addition, informal care is a widely diverse category of care. Therefore, it is assumed that care by the parents or an informal carer does not affect utility, and that any shortfall in care for the children is met by informal care or by one of the parents. However, mothers are allowed to have different preferences for leisure/home production time, which includes time to care for their child(ren), by the age of their youngest child and the number of children.

Using information on the age of children, and school and preschool hours, observed total formal childcare hours can be distributed over the children under 12 in the household, according to the number and age of children, with more care provided to preschool children than to school-age children. Potentially, this could be formally estimated using the observed information for all households, including those households with children in just one age category which may allow identification of hours of care by age of the child. To compute childcare cost, hours of care for the different children is multiplied by the relevant hourly fee which depends on age and location of the child, and then summed across the children in the family.[17] The hourly fees that families in this analysis face are determined in a similar way as in Section 4.1.

The argument in the utility function can be defined in a number of alternative ways. One option is to use total hours of childcare and divide by the number of children under 12. The preference for formal childcare could be made to depend on the number of children and/or the age of the

---

[16] More information is collected in the SIH (i.e. exact details of the type of childcare used for each child under 12, including hours of use and cost). However, this information is not made available in the basic or the expanded CURFs.

[17] Hourly fees by region of residency and age of the child can be derived from independently collected data (Department of Education, 2014).

children. Alternatively, instead of one (average) argument for childcare, the imputed childcare hours for each child in a specific age group could be used to compute average hours of care for 0-2 year old children, 3-4 year old children, 5-9 year old children and 10-12 year old children. Only families with children in the relevant age category get to make a choice for the hours of formal care. Separate preference parameters can then be estimated for formal care for children in the different age categories.

Similar to the preference parameters for income and hours of work, the preference parameters for formal childcare could be made dependent on household and parental characteristics (or even on child characteristics if those were available).

An innovation of the utility function suggested here is that it allows families to derive utility from children's participation in formal childcare, since H_{ccf} is an argument in the utility function. Not all variables and choice opportunities that are relevant to the decision-makers are observed, such as the quality of nonparental childcare (childcare options vary with regards to facilities, quality of staff, etc., as well as opening hours and fees) or the variables determining job satisfaction, apart from hours of work and wage rates for participants in market work. To capture the effects of these non-pecuniary attributes on preferences, alternative-specific stochastic error terms are introduced, which are assumed to be both independent of observed characteristics and of each other.

To simplify the choice process, the labour supply and childcare alternatives need to be organised into a finite number of predetermined categories. Regarding the link between hours of care and hours of work, some studies assume a fixed link (Ilmakunnas, 1997; Kornstad and Thoresen, 2007) between hours of market work by the mother and hours of nonparental care, with the amount of nonparental care used (at least) equalling the hours of labour supply.[18] This means that the mother cannot work and take care of her child(ren) at the same time. However, mothers out of the labour force can use formal care alternatives, or mothers can use more childcare than is required for their employment. The same is the case in the suggested approach here, where formal nonparental care is seen as a direct contributor to childcare quality, and thus to the utility function. Formal care is more than just a substitute for parental care when this is unavailable. At the same time, it is assumed that other nonmaternal care alternatives are always available if required, so that any combination of maternal labour supply and formal childcare

[18] Men are assumed not to provide any care during the working day, since they work mostly full time.

can be chosen, with any leftover time automatically filled by informal care or care by the father.[19]

The approach described here essentially involves similar steps to implement to those discussed in section 4.1, with the exception that the exogenous reduced-form estimation of childcare costs (step 1) is replaced with endogenous costs based upon the preference for formal care. Hence, step 1 involves altering the utility function that is assumed for analysis, and then estimating (unobserved) preference parameters on the basis of survey data. The approach taken to estimate these preference parameters is qualitatively similar to that of identifying preferences in a standard model of labour supply.

4.2 Modelling and simulating Activity Requirements

Activity requirements impose conditions on the receipt of certain income support payments. These conditions typically increase the personal cost of receiving affected transfer payments, and may consequently dissuade some people from choosing to take-up their eligible support payments.

A key feature of activity requirements is that they typically impose non-pecuniary costs on receipt of income support payments. In most cases, these costs are difficult to quantify and may exhibit substantial variation between affected individuals, complicating associated empirical analysis. These complications impose some non-trivial limitations on viable modelling options. Specifically, we can conceive of alternative approaches for exploring behavioural responses to previously implemented activity requirements, which could be used as the basis for projecting responses to qualitatively similar counterfactuals under a restrictive set of assumptions. We cannot, however, conceive of an approach that will project behavioural responses to all possible dimensions over which activity requirements might be altered. It is important that the reader bear this limitation in mind.

As with any decision, the influence of activity requirements can be modelled in a range of alternative ways (see Section 2.1 for discussion). We consider two alternatives here; modelling via reduced form statistical regressions (4.2.1), and structural modelling in conjunction with the labour supply decision (4.2.2).

---

[19] Although these are simplifying assumptions which may require some testing, it appears that the majority of informal care is provided at no cost or at very low cost (as observed in the 2009/2010 SIH).

4.2.1 Reduced form analysis

In this approach, the incidence of behavioural responses to activity requirements is exogenously defined. A statistical relation could be assumed that describes the probability of income support take-up. Estimating this relation on survey data observed during a given period in time would capture the impact on behaviour of activity requirements that applied during that time. This statistical description for the incidence of benefit take-up could be used in place of the standard modelling assumption of full benefits take-up.

Accounting for activity requirements through a reduced form statistical relation for take-up of transfer payments has the advantage that it is very simple to implement. It is, however, ill-suited for considering the knock-on effect that activity requirements may have on willingness to supply labour, which is one of the central justifications for such requirements. The structural approach that we discuss below can help to address this concern.

4.2.2 Structural analysis

We discussed briefly in Section 3.1.2 how activity requirements could be included in the module by associating receipt of income support payments with, for example, a time cost. We discuss this idea at further length here, noting that it would involve two substantive alterations to the assumed preference relation, relative to the "simple" module outlined in Section 3.

First, it would be necessary to include an additional 'decision' in the model. This feature would allow an individual who chooses not to work or to work few hours to decide whether to take-up income support payments that impose activity requirements, where taking up benefits would also incur a leisure cost. The tax and benefit module of the model would require a minor amendment to accommodate the decision of whether or not an individual takes up the benefits that they are eligible for at the different levels of labour supply for which their income remains below the eligibility cut-off income, but would otherwise remain unaltered.

A less obvious alteration to the preference relation that would be necessary to accommodate activity requirements in the labour supply model is allowance for a disparity between consumption and disposable income. At present, an implicit assumption of most point in time labour supply models is that consumption is equal to disposable income. Individuals who do not satisfy activity requirements, however, may be left with no disposable income. This may complicate analysis in context of some preference relations, as marginal utility often tends toward infinity as consumption approaches zero and the level of utility may not be identified at all.

Activity requirements are most likely to dissuade take-up of benefits amongst individuals who have other resources to draw upon and individuals who are only eligible for a small amount of benefit. The former subset of individuals will subsequently consume more than their disposable income during the affected time-period. A reasonable approximation of behaviour in this context will consequently require the disparity between consumption and disposable income to be taken into consideration. The easiest way to do this is to replace the assumption that consumption is equal to income with a reduced-form model of consumption as a function of a set of observed household variables, which might include disposable income, age, the numbers and age of household members, and so on.

The modelling framework described above includes two key sets of additional parameters that would require empirical identification: the reduced-form model of consumption, and the leisure costs associated with activity requirements. The reduced-form model of consumption could be estimated using standard econometric techniques and data from the Household Expenditure Survey. The leisure costs or psychological/stigma costs associated with activity requirements represent unobservable preference parameters that would need to be estimated in the same way as other preference parameters in the model.

As for childcare, a reasonable approximation of the data is likely to require that allowance be made for the disparity of individual-specific behaviour and circumstances. In the current context, this means accommodating the error term of the reduced-form regression for consumption in the evaluation of utility, and allowing for heterogeneity in the time cost associated with activity requirements. Although we are unaware of any model that has implemented this sort of structure previously, it is not qualitatively different from now-standard methods used to simulate labour supply more generally.

Hence, the approach described above provides a fairly straight-forward way of reflecting the impact of activity requirements in place at a given point in time. However as mentioned at the start of this section, it is important to recognise that the unobservable nature of the leisure costs associated with activity requirements complicates consideration of responses to associated policy counterfactuals.

5. Conclusion

The methods required for projecting labour supply responses to alternative policy contexts are now well established in the existing economic literature, and there are a number of excellent examples of practical implementations for the Australian policy context. Without on-going model development, however, the ability of existing models to accommodate evolving policy priorities becomes compromised. This is arguably already the case, as neither of the two principal microsimulation models developed for Australia (MITTS and STINMOD-B) are adapted to explore joint decisions concerning employment and childcare, or the influence of income support payments for individuals around retirement age and/or facing health issues.

This report discusses the issues involved in maintaining model capacity to ensure that it remains policy relevant, with reference to practical examples of childcare and activity requirements. Our discussion highlights the flexibility of the utility framework that is the current foundation for best-practice economic analysis of behavioural responses to policy change. It also focusses attention on the difficulties associated with parameterising models based on the utility approach, and the associated importance of model validation to ensure that output is suitable for the purpose of policy design.

With the steady advance of computing power, our capacity to test the incentives embodied by policy alternatives continues to grow. Taking advantage of this opportunity, however, requires on-going investment in modelling work, motivated in large part by the evolving policy environment and priorities of policy makers. This report helps to clarify the practical issues involved.

References

Andersson, I., Brose, P., Flood, L., Klevmarken, N. A., Olovsson, P. and Tasizan A. (1992), *MICROHUS - A microsimulation model of the Swedish household sector*. International Symposium on Economic Modelling: Gothenburg.

Basu, N., Pryor, R. and Quint, T. (1998), "ASPEN: A microsimulation model of the economy", *Computational Economics*, 12, pp. 223-241.

Bourguignon, F., Fournier, M. and Gurgand M. (1998), "Distribution, development and education: Taiwan, 1979-1994", unpublished manuscript, Delta (Joint research Unit, ENS-CNRSEHESS), Crest (Laboratoire de Microéconométrie) and Centre d'études de l'emploi and Lamia (Université Paris I).

Browning, M. Lusardi, A. (1996), "Household saving: micro theories and macro facts", *Journal of Economic Literature*, 34, pp. 1797-1855.

Buddelmeyer, H., Creedy, J. and Kalb, G. (2007), *Tax Policy Design and Behavioural Microsimulation Modelling*. Cheltenham: Edward Elgar.

Cai, L., Kalb, G., Tseng, Y.-P. and Vu, H.H. (2008), "The Effect of Financial Incentives on Labour Supply: Evidence for Sole Parents from Microsimulation and Quasi-Experimental Evaluation", *Fiscal Studies*, 29, pp. 285-325.

Caldwell, S. (1997), *Corsim 3.0 User and Technical Documentation.* Ithaca: New York.

Callan, T., O'Donoghue, C. and O'Neill, C. (1996), "Simulating welfare and income tax changes: the ESRI tax-benefit model", *Economic and Social Research Institute*.

Callan, T. and Van Soest, A. (1996) "Family labour supply and taxes in Ireland", *ESRI Working Paper No*. 78, Dublin: The Economic and Social Research Institute.

Carrol, C. (1992), "The buffer-stock theory of saving: Some macroeconomic evidence", *Brookings Papers on Economic Activity*, 23, pp. 61-156.

Creedy, J., Duncan, A.S., Scutella, R. and Harris, M. (2002) *Microsimulation Modelling of Taxation and The Labour Market: The Melbourne Institute Tax and Transfer Simulator*. Cheltenham: Edward Elgar.

Creedy, J., Hérault, N. and Kalb, G. (2009), "Abolishing the tax-free threshold in Australia: simulating alternative reforms", *Fiscal Studies*, 30, pp. 219-246.

Creedy, J., Hérault, N. and Kalb, G. (2011), "Tax Policy Design and the Role of a Tax-Free Threshold", *Public Finance and Management*, 11, pp. 338-364.

Creedy, J. and Kalb, G. (2005), “Discrete Hours Labour Supply Modelling: Specification, Estimation and Simulation”, *Journal of Economic Surveys*, 19(5), 697-734.

Creedy, J. and Kalb, G. (2005), “Measuring Welfare Changes in Labour Supply Models”, *Manchester School*, 73, pp. 664-685.

Creedy, J. and Kalb, G. (2006), *Labour Supply and Microsimulation: The Evaluation of Tax Policy Reforms*. Cheltenham: Edward Elgar.

Creedy, J., Kalb, G. and Scutella, R. (2004) Evaluating the Income Redistribution Effects of Tax Reforms in Discrete Hours Models. In Bishop, J. and Amiel, Y (eds.), *Studies on Economic Well-being: Essays in the Honor of John P Formby*. Research on Economic Inequality, 12.

Creedy, J. and van de Ven, J. (1999), “The effects of selected Australian taxes and transfers on annual and lifetime inequality”, *Australian Journal of Labour Economics*, 3, pp. 1-22.

Curry, C. (1996), “PENSIM: A dynamic simulation model of pensioners’ income”, *Government Economic Service Working Paper* 129.

Dalton, T. and Ong, R. (2007), “Welfare to work in Australia: disability income support, housing affordability and employment incentives”, *European Journal of Housing Policy*, 7, pp. 275-297.

Deaton, A. (1991), “Saving and liquidity constraints”, *Econometrica*, 59, pp. 1221-1248.

Department of Education (2014), 2013 National Early Childhood Education and Care Workforce Census, available on (accessed on 9 August 2014): http://docs.education.gov.au/system/files/doc/other/nwc_national_report_final.pdf.

Doiron, D. and Kalb, G. (2005), “Demands for Childcare and Household Labour Supply in Australia”, *Economic Record*, 81, pp. 215–236.

Dueé, M. and Rebillard, C. (2004), “Old age disability in France: A long-run projection”, *INSEE Working Paper* G2004/02.Duncan, A. and Harris, M.N. (2002), “Simulating the effect of welfare reforms among sole parents in Australia”, *Economic Record*, 78, pp. 249-63.

Falkingham, J. and Lessof, C. (1992), “Playing god: the construction of LIFEMOD”. In R. Hancock and H. Sutherland (eds), *Microsimulation Models for Public Policy Analysis: New Frontiers*. STICERD Occasional Paper, London School of Economics: London.

Fjærli, J.A.E., Gravningsmyhr, H., Holmøy, A.M.K. and Lian, B. (1995), “The Norwegian microsimulation model LOTTE: Applications to personal and corporate taxes and social security benefits”, *Microsimulation Unit Discussion Paper* MU9504.

Giannarelli, L. (1992), *An Analyst's Guide to TRIM2*. Urban Institute Press: Washington D.C.

Harding, A. (1993), *Lifetime Income Distribution and Redistribution: Applications of a Microsimulation Model*. North-Holland: London.

Hérault, N., Kalb, G. and van de Ven, J. (2014), "The Effects of Income Support Settings on Incentives to Work", *Melbourne Institute Report* commissioned by the Australian Government Department of Education, Employment and Workplace Relations (DEEWR).

Hoynes, H.W. (1996), "Welfare transfers in two-parent families: labor supply and welfare participation under AFDC-UP", *Econometrica*, **64**(2), 295-332.

Ilmakunnas, S. (1997), "Female Labour Supply and Work Incentives", Labour Institute for Economic Research, Finland.

Kalb, G. (2000), "Labour supply and welfare participation in Australian two-adult households: Accounting for involuntary unemployment and the 'cost' of part-time work", Working Paper No. BP-35, Centre of Policy Studies and the Impact Project, Monash University.

Kalb, G. (2004), "The Use of Simulation Models in Policy Analysis", *Australian Journal of Labour Economics*, 7, pp. 1-12.

Kalb, G. and Lee, W.-S. (2008), "Childcare Use and Parents' Labour Supply in Australia", *Australian Economic Papers*, 47, pp. 272-295.

Kalb, G. and Thoresen, T.O. (2009), "Behavioural Microsimulation: Labour Supply and Childcare Use Responses in Australia and Norway" in *New Frontiers in Microsimulation Modelling*, eds.: A. Harding, P. Williamson and A. Zaidi, 231-261. Surrey: Ashgate.

Keane, M. and Moffitt, R. (1998), "A structural model of multiple welfare program participation and labor supply", *International Economic Review*, 39(3), 553-589.

Kimball1, M.S. (1990), "Precautionary saving in the small and in the large", *Econometrica*, 58, pp. 53-73.

King, A., Bækgaard, H. and Robinson, M. (1999), "The base data for DYNAMOD-2", *NATSEM Technical Paper* 20.

Kornstad, T. and Thoresen, T.O. (2007), "A Discrete Choice Model for Labor Supply and Child Care", *Journal of Population Economics*, 20, pp. 781–804.

Lucchino, P. and van de Ven, J. (2013), "Modelling the Dynamic Effects of Transfer Policy: The LINDA Policy Analysis Tool", *Melbourne Institute Working Paper Series* 20/13.

Lucas, Robert (1976), Econometric Policy Evaluation: A Critique. In Brunner, K. and Meltzer, A. *The Phillips Curve and Labor Markets*. Carnegie-Rochester Conference Series on Public Policy 1. New York: American Elsevier

Maddala, G.S. (1983) *Limited Dependent and Qualitative Variables in Econometrics*. New York: Cambridge University Press.

Moffitt, R. (1983), "An economic model of welfare stigma", *American Economic Review*, **73**(2), 1023-35.

Nelissen, J. H. M. (1994), *Income Redistribution and Social Security: An Application of Microsimulation*. Chapman and Hall: London.

Orcutt, G. (1957), "A new type of socio-economic system", *Review of Economics and Statistics*, 58, pp. 773-797.

Orcutt, G., Caldwell, S. B. and Wertheimer, R. (1976), *Policy Exploration Through Microanalytic Simulation*. Urban Institute: Washington D.C.

Redmond, G., Sutherland, H. and Wilson, M. (1998), *The Arithmetic of Tax and Social Security Reform: A Users' Guide to Microsimulation Methods and Analysis*. Cambridge University Press: Cambridge.

Sutherland, H. (2001), "EUROMOD: An integrated European benefit-tax model", *EUROMOD Working Paper* EM9/01 Final Report.

Zaidi, A. (2007), "The SAGE model: A dynamic microsimulation population model for Britain". In W. Barnett (eds), *Modelling Our Future: Population Ageing, Health and Aged Care*. Emerald Group Publishing Limited: London.