

Australians over the coming 50 years: providing useful projections

Simon Kelly¹, Anthony King
National Centre for Social and Economic Modelling
Canberra, Australia

June 2001

Abstract

Dynamic microsimulation models have long promised to provide a new level of detail on the image of the future. The Australian model DYNAMOD is one such example. DYNAMOD models behaviour at the individual level, has used groundbreaking techniques to achieve this, and can produce an incredibly in-depth view of Australia's population out to 2050. However, it has proved to be a long process getting DYNAMOD to the point where it is seen as a valuable tool by potential users – much longer than originally anticipated. The main reason for this has been the lack of attention to alignment in the original model design, and the work subsequently entailed in adding alignment to the model. Alignment is important both in order to enable the specification of future alternative scenarios and to match model output with actual outcomes – such as, labour force outcomes to date. Recognition of the importance of alignment as an essential element for a useful model has seen an emphasis on these aspects of development in recent versions of DYNAMOD.

This paper describes some of the alignment processes that have been introduced into the most recent versions of DYNAMOD, with a focus on the techniques used to align processes which are modelled using survival functions. Examples, such as the alignment of fertility, are used to highlight the incremental improvements that have been made to the model. The examples illustrate the benefits of moving from the original unconstrained survival function model to a model using adjusted survival functions and finally to a model which uses survival functions as part of a transition process. Only now, with this capacity for alignment in place, can DYNAMOD start to claim the plausibility, flexibility, transparency and user utility to serve as a useful model.

¹ Address for correspondence: NATSEM, University of Canberra, 170 Hayden Drive, Bruce ACT 2601, Australia.
Email: simonk@natsem.canberra.edu.au
Website: www.natsem.canberra.edu.au

1. Introduction

Dynamic microsimulation models have long promised to provide an image of the future with a level of detail not possible using other projection techniques. The good news is that some of these models have begun fulfilling this promise – they are operational now. The US Social Security Administration has adopted the Cornell University dynamic microsimulation model, CORSIM, as one of its tools in support of social security policy development (Morrison, 2000). While DYNACAN, the Human Resources Development Canada dynamic microsimulation model is an integral part of the policy development process for the Canada Pension Plan (Morrison, 2000). These examples show that dynamic microsimulation can move from “academic exercises” to implementation – where the tool is being used in real world problem solving and policy development.

For this transition from prototype to operational model, it is important that the developer, users and decision makers are all convinced that the answers coming from these models are reliable. The model projections must be credible. Sargent (2000) suggests model verification and validation are the starting points to gaining this credibility. Verification ensures the internal workings of the model are correct while validation ensures the output is reasonable. King (1995) provided four guiding criteria for development, which assist verification and validation – plausibility, flexibility, transparency, and reflection of user needs. Achievement of these criteria is seen to increase the understanding and utility of the model and, with repeated use, confidence and acceptance of the results from the model grows. Credibility and usefulness will then have been achieved.

The early work of NATSEM was concentrated on the development of two core microsimulation models - STINMOD, a static microsimulation model designed primarily for tax-transfer analysis; and DYNAMOD, a dynamic microsimulation model of the Australian population. STINMOD quickly gained acceptance, is used by a number of key government agencies, and is regularly updated. In contrast, it has proved to be a long process getting DYNAMOD to the point where it is seen as a valuable tool by potential users – much longer than originally anticipated. A general reason for the different fortunes of STINMOD and DYNAMOD is the far greater complexity of dynamic modelling. This complexity has meant that development, verification and validation of DYNAMOD has been a much more involved task than it was for STINMOD.

So, perhaps the slow development of DYNAMOD is unremarkable. However, there is a particular aspect in the history of DYNAMOD that has been recognised as initially hampering, and then delaying, its shift to the status of a useful model. This is the degree to which the model includes the capacity for alignment. Alignment is important both in order to enable the specification of future alternative scenarios and to match model output with actual outcomes – such as, labour force outcomes to date.

Very little attention was given to alignment in the original DYNAMOD model design and a number of papers questioned whether DYNAMOD could become a credible and useful model in its then form (see Hellwig, 1995; King and Bækgaard, 1997; or Bækgaard, 1997). The key element missing was identified to be alignment processes; a point supported by the experiences of the successful CORSIM and DYNACAN dynamic microsimulation models mentioned above. Recognition of the importance of alignment as an essential element for a

useful model has subsequently seen an emphasis on these aspects of development in recent versions of DYNAMOD.

Why is model alignment important? Alignment serves two broad purposes:

1. Correcting for imperfect estimation; and
2. Scenario setting.

Mechanisms to achieve these purposes are variously referred to as benchmarking, calibration or alignment – we use the term alignment to refer to the whole range of techniques. Sometimes it is alignment of model outcomes to actual aggregate data (commonly referred to as benchmarking), or adjustment of simulation equations to yield results that tally with observed outcomes (often referred to as calibration), or adjustment so model outcomes match user-specified aggregate scenarios. Use of alignment for scenario setting is a natural and uncontroversial application. There are, however, two schools of thought on the use of alignment to correct for imperfect estimation. Some would argue that imposing alignment compromises the integrity of the model, by over-riding the estimated behavioural equations. On the other hand, we never have perfect estimation data – no matter how good are our estimation techniques – and there will always be the scope and the need to repair consequences of imperfect estimation. Development of DYNAMOD, at least over the course of DYNAMOD-2 and DYNAMOD-3, has reflected the second of these schools of thought. Further discussion on the role of alignment can be found in Bækgaard (2001).

The methods used to add alignment capabilities to DYNAMOD are the subject of this article, with an emphasis on the particularly challenging task of aligning processes that are modelled using survival functions. The account distinguishes three phases in the development of DYNAMOD – phases that are characterised by the degree of attention to alignment.

1. DYNAMOD-1 (1993-94): Development of the original version of DYNAMOD focused on the incorporation of techniques – survival functions – to represent individual behaviour more fully than was possible using the conventional transition probability approach. In this development environment, alignment of the output from the model with user-defined scenarios or with the real world was, at best, a secondary consideration. This version is discussed in section 2.
2. DYNAMOD-2 (1995-98): The focus of work on the second version of DYNAMOD (section 3) reflected recognition of the importance of alignment. It involved adding alignment capabilities to much of the model, including first efforts with the alignment of the processes using survival functions. New elements were also added to the model with alignment included as basic elements of these.
3. DYNAMOD-3 (1999 -): The current version of DYNAMOD (section 4) has seen considerable increase in the sophistication and coverage of the alignment of the survival function processes. The model is now at the stage where it can start to claim the plausibility, flexibility, transparency and user utility required if it is to serve as a generally useful model.

The gains in model performance, which have been achieved by the work on alignment, are illustrated by presenting the results from the modelling of fertility (a process modelled using survival functions) under the three versions.

The article concludes with a discussion of other developments with DYNAMOD, such as improved verification and documentation, and some examples of how the aggregate model performance matches benchmark data. The improvements in these areas should help to improve understanding, transparency, and credibility of the model and, thereby, usefulness.

2. DYNAMOD-1 (1993-94) – unconstrained projection

2.1 An overview of DYNAMOD-1

Descriptions of DYNAMOD-1 can be found in Antcliff (1993) and in the papers compiled in Antcliff *et al* (1996). Development of the model began in 1993, with one of the original developers having stated that the original aim was to:

‘... reflect the complexity of real life as far as possible within the model .
In particular, we were looking for a greater degree of interaction between the different processes being modelled and the use of a continuous time frame rather than annual snapshot.’ (Gruskin 1994, pp.8)

The potential power of such a model was obvious - it would provide a significant advance in the tools available for policy analysis by delivering detailed individual projections that took considerable account of the dynamic nature of the socio-economic characteristics of people’s lifetimes. It was the first model of its type in Australia and one of the first in the world to use a continuous time frame. Survival functions were seen as the way to model continuous time and the very detailed interdependencies of individual histories, while not producing significant costs in terms of computing requirements.

In dynamic microsimulation models, the ‘ageing’ of the population, in terms of the changing characteristics of individuals over time, has conventionally been modelled using transition probabilities. For every period that an individual is in scope for a particular transition, they are assigned the probabilities of each possible outcome of the transition, including the event that nothing occurs. A random number then determines the outcome. A simple case in point is a transition probability model for mortality where the probability of death in a time period is specified, typically given the age and sex of the individual. For example, let the probability of death for a given individual during a year be 0.03 (3 per cent). A random number (between 0 and 1) is then drawn and if the random number has a value less than 0.03 then the outcome for this person is death. If the random number is not less than 0.03 then the person continues on living for that year. In general, this calculation is done for every person and every period covered by the model. Moreover, every transition in a period typically occurs at the same time. In the case above everyone who is going to die in that year, dies on the first day of the year. On this same day, all marriages happen, all births happen, all entries into employment happen, and so forth.

In contrast to this approach, DYNAMOD uses survival functions to model some events. Survival functions model the time until an event occurs, that is a prediction from a survival model is a future point in time when the event is expected to occur. Survival functions generate a ‘failure date’ for a particular event. The date is a function of the underlying hazard and the characteristics of the individual at the time of calculation. In modelling this function,

the underlying hazard is represented by a baseline value and the influence of various characteristics by a vector of covariates. The general model is:

$$h_t = h_{0t} + \sum \beta_{it} z_{it}$$

where h_t is the chance of experiencing the event at time t ; h_{0t} is the baseline hazard at time t ; z_{it} is the observed value of i th covariate at time t ; and β_{it} is the regression parameter associated with the i th covariate at time t . The date of the event is then calculated using h_t and a random number.

$$d = t + \text{int} \left(\frac{-\ln(r)}{\exp(h_t)} \right)$$

where d is date of the future event (in months); t is the time (the current month); r is a random number drawn from a uniform [0,1]-distribution; and h_t is the chance of experiencing the event. A straightforward example of the use of survival functions can be seen in estimating the time until divorce after separation of a couple. The survival function for this event is²:

$$h_t = -5.059 + dur_1 * 2.058$$

where $dur_1 = 1$ if the number of years since the separation is equal to 1 otherwise $dur_1 = 0$. If we assume at time t a person has just separated and that a random number of 0.5 is drawn then $h_t = -5.059$ and d (the date of divorce) is $t + 109$ months.

There are two main features distinguishing survival function models from transition models. Firstly, the modelling can more easily approach continuous time by enabling events to happen at any time rather than all at the start of the period. In DYNAMOD, the survival functions operate with a period of one month (pseudo-continuous time, compared to annual models). A transition probability approach could of course also be set up to operate on a monthly time period though its specification and operation would be considerably more unwieldy. Secondly, in survival functions the variable being modelled is the length of time before an event occurs. If we consider the modelling of mortality again using a survival function, the calculation will give an estimate of time until death. A major advantage of survival functions is that calculations do not need to be repeated every period – only when there is a change in one of the background variables used in estimation of the survival function. In the case of DYNAMOD, the only characteristics that influence the age at death are: year of birth, sex and a person's disability status. As only disability status can change and rarely does, most people only have their date of death calculated once during their modelled life. For continuous time, the result is that a calculation needs to be done only once or twice rather than around 840 times (70years x 12months) using transition probabilities.

The survival function described above is a simple one. In contrast, most of the survival functions in DYNAMOD are complex and relate a transition to a range of background variables. For example, the survival function for first marriage includes references to age, cohabitation, employment status, pre-marital pregnancy, educational status and attainment. This list points to one of the attractions of survival functions – the ease of incorporating a range of characteristics as influences on a transition. Survival functions provide an elegant

² This is a simplified version of the actual equation used in DYNAMOD. The full version is found in Antcliff *et al* 1996 pp62.

and comprehensive summary of the influences in individual behaviour. To do the same using transition probabilities is awkward and cumbersome in comparison.

In DYNAMOD mortality, fertility, disability, couple formation, couple dissolution, marriage and divorce are all modelled using survival functions. The absence of suitable longitudinal estimation data precluded use of survival functions in other areas of the model, such as labour dynamics.

In summary, survival functions have three major attractions:

- they can model (pseudo) continuous time,
- calculations only need to be done when circumstances change, and
- the influence of a range of characteristics can be incorporated into a transition.

Against this list of strengths, the use of survival functions in DYNAMOD also has drawbacks. The first is that many of the functions were developed using data from a longitudinal survey undertaken in 1986. The survey collected detailed historical information on women (aged 20-59 at the time of the survey) and their partners. For the purposes of estimating the DYNAMOD survival functions, only the data for the period from 1971 to 1986 were used and, in the case of divorce, only the data from 1977 onwards. The survival functions developed therefore represent the behaviour of the relevant cohorts over the 1970s and first half of the 1980s. They may or may not well represent the behaviour of other cohorts or even these cohorts at a different time. The base year for DYNAMOD is 1986 and, as a projection moves into the 21st century, the estimated survival functions may diverge increasingly from actual or likely behaviour.

A second drawback of survival functions is that complexity makes modelling of alternative scenarios very difficult. Consider the influences on first marriage mentioned above (age, cohabitation, employment status, pre-marital pregnancy, educational status and attainment). A plausible scenario to model might be that the average age of first marriage for women changes from the current 26.2 years (24.0 in 1986) to 30.0 years. How should the coefficients on each background variable change to accommodate this scenario? The answer is that there is no straightforward manner. In this regard, the price of complexity is inflexibility.

The third area of concern with survival functions, arising from this inflexibility, is the difficulty of aligning the model processes with the real world. The need for alignment arises whenever the estimation data is imperfect. In the case of the survival functions used to model demographics in DYNAMOD, their estimation is based on behaviour prior to 1986. As DYNAMOD starts in 1986, we have 15 years of actual (albeit aggregate) data on more recent behaviour to compare with the picture generated by DYNAMOD over this period. If these more recent data can be taken into account, then DYNAMOD would be on a much stronger footing for projecting into the future. Addressing this problem is similar to the one above and there is no simple straightforward way of aligning the data or adjusting the survival functions to reflect changed behaviour.

Fourthly, when considering group behaviour, with the exception of life expectancies in the case of mortality, people tend not to think in terms of survival functions. Instead, the more familiar language is in terms of collective rates, such as fertility rates or divorce rates. These are actually transition probabilities that can be modelled directly and explained in a meaningful manner. In contrast, survival functions are more complex and cannot be described in such a straightforward way. For this reason, people are more likely to understand and be

comfortable with transition probabilities. They can comprehend and think in terms of fertility rates and divorce rates but find it difficult to translate these concepts into survival terms. With the exception of life expectancies noted above, there are no equivalent or common examples for survival. This means that a model based on survival functions is less transparent and has to work much harder to gain broad credibility than one based on transition probabilities.

Finally, we return to the argued efficiency of survival functions in terms of the number of times calculations need to be performed during a model run. The functions only need to be recalculated when background variables change, but since many of the DYNAMOD survival functions include a range of background variables subject to change over the course of a simulation, survival functions for an individual may need to be recalculated many times. The efficiency argument for survival functions is not as great as it may seem at first glance.

In summary, the survival functions in DYNAMOD have had the following drawbacks:

- the functions are largely based on longitudinal data from a 10-15 year period ending in 1986 which may not well represent current or future behaviour;
- the form of the functions make it difficult to model different scenarios;
- the alignment to actual data is difficult;
- the functions can be complex and hard to comprehend; and
- the efficiency advantages are not as clear cut as initially thought.

Against the conceptual attractions of survival functions, there are disadvantages – many of which relate to the capacity for alignment. Considering the prototyping nature of DYNAMOD-1, the complexity of the functions and the data used to construct the functions, it is perhaps unreasonable to expect the issue of alignment to have received much attention. That said, the lack of this attention handicapped the model in its quest for acceptance. The capacity for alignment was directly built into only one of the processes in DYNAMOD-1 – mortality. This appears to have stemmed from the way in which mortality is handled rather than through an explicit concern with the need for alignment. With mortality, the model operates with transition probabilities as input data (which can be readily altered). These are translated into survival function terms within the model. Otherwise, alignment was built into neither the other survival function processes (fertility, disability, couple formation, couple dissolution, marriage and divorce), nor the processes modelled using transition probabilities (education and labour force activity).

The gap that can exist between the output of an unaligned dynamic microsimulation model and that which would inspire the confidence needed for using the model in real world applications is illustrated below with the example of the modelling of fertility in DYNAMOD-1. The example shows the power of survival functions while also illustrating the drawbacks.

2.2 Unaligned fertility in DYNAMOD-1

The original method used to model fertility in DYNAMOD divides births into three types: premarital, first marital, and subsequent marital births. There is a pregnancy event associated with each of the birth types and survival functions are used to predict the time until the possible occurrence of pregnancy among women who are ‘in scope’. At any time, a woman can only be in scope for one type of birth. Every woman aged 15-35 who has not already given birth and has never been in a marital union (*de jure* or *de facto*) is considered in scope

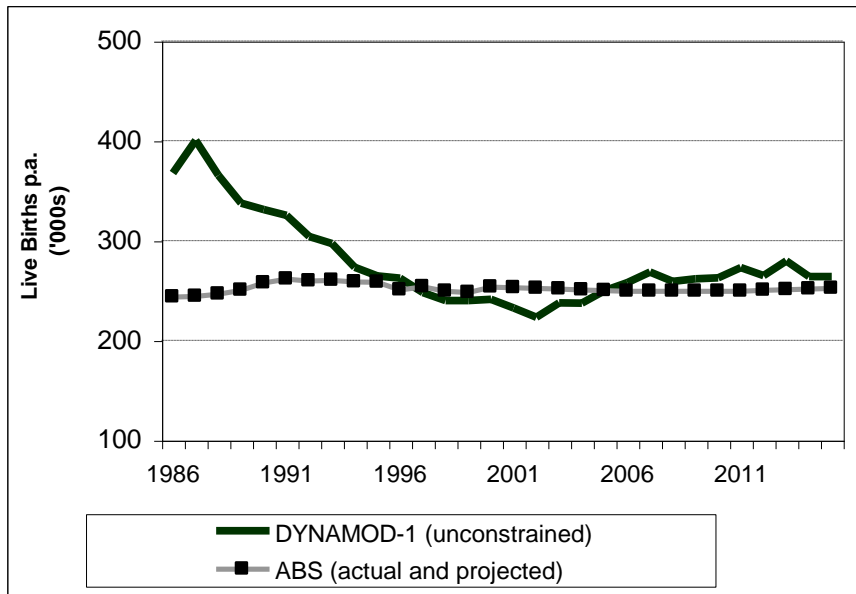
for a premarital birth. If a woman is in scope for a premarital birth then the time until pregnancy is influenced by age, educational participation and qualifications, and employment status. Women in scope for first marital births are aged under 50, have not been married for more than 20 years and have never given birth. In addition to the premarital influences on time until pregnancy, pregnancy for first marital birth is influenced by current marital status, employment status of the husband and duration of the marital relationship. Those in scope for second or subsequent birth are restricted to those aged under 50, married less than 20 years and, naturally, have already given birth. In addition to the premarital influences, the second or subsequent birth is influenced by the number of previous births, time since the last birth and duration of full-time employment. The resulting duration until the pregnancy start date coming from the survival function can be any time from now until 'eternity'. Eternity is defined as any time after the end of the simulation. A start date of eternity effectively means that a pregnancy will not occur. The actual specifications for these survival functions can be found in Bracher (1993).

Each woman predicted to give birth has a pregnancy start date placed in a data structure known as the 'crystal ball'. The crystal ball is checked each month and if a pregnancy start date equals the current month then the pregnancy begins. Until the pregnancy start date is reached, any change in the influences will result in a recalculation of the date. However, once pregnant, the birth cannot be superseded by any other event except the death of the mother. When the birth does occur, further calculations are done to assign the sex of the child (or children, if a multiple birth is simulated).

The detailed description above shows some of the power of survival functions. The prediction of when a woman will have a child is based on her circumstances at that time. Each of these many circumstances has a weighted influence on the start date of the pregnancy. Once the calculation is done, if the circumstances don't change then the calculation is not done again. If the influences do change, then the calculation is redone considering the changed circumstances. It is an efficient, comprehensive and elegant solution. At this individual level, the function follows a natural estimation process. If a woman is 17 years old and still at school, pregnancy is a certain amount of time away. If two years later she marries then the time until pregnancy needs to be recalculated and will probably be found to have come forward. If the couple separate before she is pregnant, then a third pregnancy start date will be calculated; if she is already pregnant when separation occurs, then the birth will go ahead. At this level, the processes are logical, easy to describe, easy to understand and bring together all the relevant significant influences. Without continuous time and survival functions, this level of detail would be extremely difficult to model.

Against these advantages is the issue of alignment. Figure 1 shows the number of births predicted by DYNAMOD with unaligned fertility compared with the actual number of births reported by the Australian Bureau of Statistics (ABS) until 1998 and the ABS projection of future births (under the Series II scenario where fertility falls to 1.75 births per woman by 2006) from 1999 on.

Figure 1: Live Births 1986 - 2016



Sources: ABS Births Australia 1998 Cat no. 3301.0; Population Projections 1997-2051, Cat no. 3222.0 Table 4.5.

Figure 1 shows the model outcomes diverging substantially from the benchmark data on actual births for the first few years of the simulation, with the simulation being about 50 percent higher over the period 1986-1989. The reason for this divergence relates to incomplete data for individuals in the simulation base data. Where a start date is not known for a particular characteristic, an estimate is made using other information. If this technique cannot be used, a random date or the simulation start date is assigned. The result of this arbitrary assignment is that some event influences are ineffective at the start of the simulation. For example, duration of full-time employment reduces the chance of second or subsequent childbirth. In DYNAMOD, duration of full-time employment will be zero at the start of the simulation and the influence of this characteristic will be missing for the first few years of the simulation. Over time, the model will stabilise and come in to line with ABS figures (as seen in Figure 1). However, the damage has been done. The unusually high number of births over 1986-1989 result in this age group being particularly large through successive years of the simulation. The divergence is felt not only in the population profile but also through interaction with a number of the simulated outcomes. An obvious example is the impact on the projected numbers of school students, and then tertiary students in coming years. These are thrown out by the large number of projected births. A number of other related outcomes will also be biased by the large birth cohort.

Besides being unsatisfactory for the modeller, model outcomes, such as that shown in Figure 1 provide no basis for establishing model credibility. Potential users could justifiably suggest that if the results diverge from known benchmarks by up to 50 per cent in the first few years – what use are the future projections? User credibility requires that the model, at the very least, can produce outcomes that line up with actual aggregate benchmark data.

Overall, the original version of DYNAMOD showed that the use of survival functions were both an elegant and viable option for dynamic microsimulation modelling of a population – in technical terms. In practical terms, however, it also highlighted that more work needed to be done to align the summation of this individual behaviour to aggregate benchmarks (based on actual data or constructed future scenarios), if the model was to prove useful.

3. DYNAMOD-2 (1995-98) – adding alignment

3.1 An overview of DYNAMOD-2

Development of DYNAMOD-2 (King, Bækgaard and Robinson 2000) was a period of consolidation. Some new modules were added to the model – notably earnings and overseas migration – though the focus of the effort was on adding alignment to the model. While the emphasis with DYNAMOD-1 had been on technical design, the emphasis of DYNAMOD-2 was more pragmatic – to make the model useful. This aim was to be achieved by addressing the shortcomings of the original ‘unconstrained’ version.

A variety of forms of alignment were introduced in DYNAMOD-2:

- An alignment capacity was added to the labour force processes, constraining the labour force transitions to match user-input time series of unemployment rates and labour force participation rates (Bækgaard 2001). This alignment operates by using the combination of the transition probabilities and random numbers each period to rank individuals according to their simulated likelihood of making a transition (say, from full-time employment to unemployment) and then selecting cases in rank order until the exogenous target (for example, the unemployment rate) is met.
- Alignment of education processes was introduced in a similar fashion, with the annual transitions constrained to match user-specified key aggregate education transition rates, such as the rate of progression from secondary schooling to university (Robinson and Bækgaard 2001).
- DYNAMOD-2 saw completion of the earnings module, including the provision for alignment (Bækgaard 1996). Individual changes in earnings are set to match exogenous aggregate earnings growth rates (actual benchmarks or scenario targets) for groups of workers defined by characteristics such as employment status, sex and skill category.
- Other new additions with DYNAMOD-2 were the overseas migration modules (King, Walker and Bækgaard 2001). In the cases of both immigration and emigration, the numbers and characteristics of migrants are governed by selection against detailed targets.

The combined effect of the work described above was the addition of considerable flexibility to DYNAMOD with the facility for matching exogenous targets introduced across a range of areas. This then provided the capacity for aligning the model outcomes to match actual outcomes over the period from the 1986 base year to the present, and also to set alternative future scenarios for the model projections. The alignment mechanisms, however, still did not cover the whole model. It will have been noticed that alignment of the demographic processes based on survival functions – the example used above as demonstrating the importance of alignment – is missing from the above list of developments.

Alignment of the demographic processes modelled using survival functions posed the greatest challenge. The relative merits of survival functions and transition probabilities remained a major issue with particular regard to the ease of alignment (Bækgaard and King, 1995; Galler 1997). Finally, however, a possible solution was identified for the approximate alignment of

continuous-time processes (Bækgaard 2001) and implemented in the fertility module of DYNAMOD-2. This alignment is described below.

3.2 Approximately aligned fertility in DYNAMOD-2

We saw in Figure 1 that unaligned modelling of fertility in DYNAMOD-1 had produced less than satisfactory results, with ramifications for many areas of the model outcomes. In response, a method was developed in DYNAMOD-2 to achieve an approximate alignment of fertility. This entailed making adjustments to the survival function by adding an adjustment factor that was calculated period-by-period from a comparison of unaligned results for that period with benchmark data for the same period. The method is described in more detail below.

The fertility survival functions consist of a constant (called the baseline value) and a number of coefficients reflecting the influences of the background variables (see section 2.2). The alignment procedure used in DYNAMOD-2 made a period-by-period adjustment to this baseline value. At the start of each year, an adjustment figure was calculated and added to the baseline value. Referring back to the general form of a survival function:

$$h_t = h_{0t} + \sum \beta_{it} z_{it}$$

h_{0t} was no longer a constant but rather changed on an annual basis, that is

$$h_{0t} = h_0 + \delta_T$$

where h_0 is the original baseline value; and δ_T is the calculated adjustment for year T .

The adjustment was calculated by undertaking an ‘indicative’ unaligned run at the start of the period. This unaligned run used the original baseline value. The unaligned numbers of births forecast for the upcoming year were compared with benchmark data. The difference between the unaligned and benchmark values was used to calculate a correction to the original baseline value. The survival function incorporating the adjusted baseline value was then used to model fertility for that period.

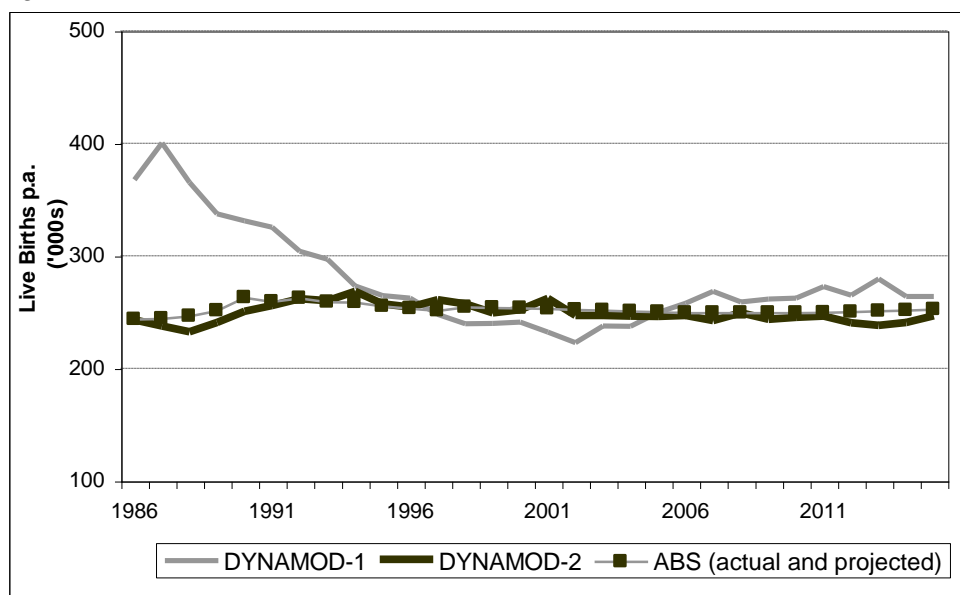
While the adjustment could be calculated every period (i.e. each month), this was considered impractical and an annual basis was adopted. Under DYNAMOD-1, the timing of a birth was only re-evaluated when one of the background variables changed. Now, each person was evaluated one additional time per year to ensure the model matched the required benchmark outcomes.

The benchmark targets for births were available both as an aggregate number for the year and disaggregated by the mother’s age group. The latter of these was used as the alignment data in DYNAMOD-2. It was felt that benchmarking at this level would improve the model results in terms of age-specific fertility numbers. The task of aligning by age within the model was not simple as the survival functions were based on type of birth (premarital, first marital, and subsequent marital) rather than age. The solution required adjustments to be calculated by age group of the mother and then converted to three ‘type of birth’ weighted average adjustments - one for the each of the baseline values used in the survival functions.

The above approach did not guarantee an exact fit with aggregate benchmark data for a number of reasons. Firstly, the alignment of births by mother's age group assumed the age distribution of females in the model matched the benchmark (ABS) age groups. This was not necessarily the case due, for example, to sample selection (the simulation uses a one percent sample of the 1986 population) and Monte Carlo variability in the simulation of death, migration and fertility. Secondly, the alignment values were only calculated once per year and did not take into account changes that may have occurred throughout the remainder of the year. A change in circumstances throughout the year could affect the start date for pregnancy and therefore change the number of births. Finally, the use of weighted average adjustments for baseline value of birth survival function may not give the correct overall tally. The number of births that occur in the forthcoming year depends on the number with a time until the birth of less than 12 months. This time is based on an exponential function of the baseline value and its influences. This means that a change of x in the baseline value will produce a change of e^x in the duration. The result is that a small change in the weighted average adjustment to the baseline value will result in a large change to the time until the births occur. Clearly, a small error in this value could result in a large increase or decrease in number of births planned to occur in the forthcoming year.

Despite these limitations, the results from incorporation of the approximate alignment procedure showed a very considerable improvement over the unaligned model outcomes (Figure 2). The divergence in the early years had been removed and the feedback through annual adjustments saw the simulation track the benchmark quite closely. The cost in terms of model efficiency had been that females in scope for a possible birth were processed one additional time during the year. This was considered a very small cost for the improvement.

Figure 2: Actual and Simulated Live Births 1986 - 2016



Sources: ABS Births Australia 1998 Cat no. 3301.0; Population Projections 1997-2051, Cat no. 3222.0 Table 4.5.

The DYNAMOD-2 fertility results, while a significant improvement over the original model, still had their limitations. Firstly, as noted above, the alignment implicitly assumed the aggregate population and age distribution for each period aligned with actual data. The model aligned the simulated number of births with the benchmark data using aggregates rather than rates. For example, the alignment process attempted to produce around 250,000

births per year rather than, say, 16 births per 1000 people. Even if only a third of the population were used, there would still be 250,000 births per annum. While this alignment would be suitable for many applications, it was unsuitable for user scenarios involving changing population structures. Model testing and development were also hindered, as a smaller “test” population could not be used. Secondly, examination of the alignment results by mother’s age group showed outcomes which were not as good as those for the overall alignment. Remembering that the survival functions operate on type of birth (pre-marital, first marital and subsequent), undertaking the alignment by mother’s age group meant that the adjusted baseline value for the type of birth was a weighted mean of the age group values. Experience with the model revealed that the use of this average value biased the output towards certain types of birth. In particular, pre-marital births tended to be too high. A consequence of this was that while aggregate births were aligned, the mother’s age at first childbirth was not increasing over time, as it should have been if it was to match other benchmark data. These limitations were addressed in DYNAMOD-3.

4. DYNAMOD-3 (1999-) – further development

4.1 An overview of DYNAMOD-3

DYNAMOD-3 is the next evolutionary step in the development of NATSEM’s dynamic microsimulation model. The focus is now clearly on increasing user acceptance of the model. Three specific areas have been addressed to improve acceptance – the platform, additional modules and more sophisticated alignment. In addition, considerable effort has been expended on increasing the reliability of the code by thoroughly testing the code for tolerance of exceptional cases and programming faults. DYNAMOD-3 is now delivering the usefulness users require. The results are reliable, plausible, flexible, and transparent.

Perhaps the biggest step in the development of DYNAMOD-3 was the shift to a personal computer (PC) platform. The move, from a model that required a high performance Unix-based computer to a model that runs on a desktop computer overcomes a major obstacle for users. Most potential users do not have easy access to a Unix based computer and therefore use of DYNAMOD required the purchase of a specialist simulation computer. With the move to a PC, anyone with a reasonably powerful desktop computer can simply load the software and run the simulation. A second advantage of the move is from a development perspective. The popularity of the PC environment means that the range of development tools available is greatly increased. These tools allow development that is more integrated and permit early identification of coding problems thereby speeding development and improving the quality. In addition, these tools have enabled a number of persistent, minor but irritating bugs to be identified and remedied. The third and unexpected benefit of the move was the reduction in simulation time. The run time on the PC version has been reduced by around 40% over the Unix version.

Modules covering personal income taxation, household assets and superannuation have been added. Earnings, which can now be computed in either nominal or real terms, have income tax calculated and recorded. The taxation rates used within the simulation can be either actual historical and current rates or user-defined rates. This is the type of flexibility missing in earlier versions. Now the model can produce outcomes based on gross or net income, actual or user-defined taxation rates, real or nominal dollars. In the new assets module, household assets such as a family home; cash in the bank; rental property investments; stocks and shares

are accumulated and disposed of during people's lives (Kelly, 2000a). The value of these assets and any associated debts are tracked, evaluated and updated throughout the simulation. The ability to model of wealth accumulation is viewed as major advance for DYNAMOD. This distribution of wealth information will further increase the range of research areas for which the model is suited. The final new module is one to capture the accumulation of superannuation or retirement funds (Kelly, Percival and Harding 2001). Again, the module has been designed to utilise either historical benchmark data or user-defined input. The ability to allow the user to set, say the future earnings rates for retirement funds, increases the flexibility and usefulness of the model. With the cost of supporting the increasing retired population on the political agenda of most countries including Australia, one area where DYNAMOD is particularly useful is in projecting the financial position of retirees under a range of scenarios.

4.2 Alignment in DYNAMOD-3

Alignment, at best an afterthought in DYNAMOD-1, is a feature of each of the new modules in DYNAMOD-3 and the alignment processes have been improved in a number of other modules. A number of areas of the model have had alignment introduced or strengthened in DYNAMOD-3:

- Improved labour force alignment. In response to poor simulation outcomes at higher levels of disaggregation for age groups in the previous versions, DYNAMOD-3 reduced the size of the age groups (i.e. increased the number of age groups) to improve the alignment. Labour force transitions are now aligned by 10-year age group (15-24, 25-34, 35-44, 45-54, 55+) for males and females. More recent benchmark data is now used to better reflect current transition rates and changes were made to the process of people transitioning between full-time and part-time employment. Again these changes improved disaggregated alignment. There are significant code changes to the model to facilitate this.
- Improved marriage alignment. In version 2 of the model, simulation outcomes in relation to proportion by marital status matched benchmark data when viewed in cross-section. However, longitudinal evaluation of couple formation and dissolution showed that the rates were too high. In DYNAMOD-3, a time-based adjustment factor was added to the baseline value of the first marriage survival function. At time t an adjustment δ_t was added to the original baseline value h_0 to produce a corrected baseline value for that period h_t :

$$h_t = h_0 + \delta_t$$

The value of δ_t is calculated by comparing the simulated aggregate number of births in a specific year with benchmark data. Use of this adjustment corrected the misalignment and calibrated the simulation outcomes.

- More improvements were made to the alignment of fertility with disaggregated benchmark data. This example is discussed further below.
- The personal wealth module integrated alignment into the ownership decision. The decision to purchase shares, a family home, rental investment properties were all influenced by benchmark alignment data.
- Superannuation coverage and account balances are initialised with a snapshot in 1993). This cross-sectional snapshot is aligned with benchmark data.

The alignment of probability transitions such as those used in the wealth module are normally done in the following manner. The current simulated number in one state is compared with an exogenous target number. The difference is then expressed as a proportion and is used as the likelihood of making a transition. Each individual within scope to make this transition is assigned a random number. Selection is made based on the random number and transition probability. An example of the use of this technique is in deciding whether to invest shares i.e. make the transition from non-shareowner to shareowner. The simulated proportion of shareowners of a particular sex, age and income is compared with benchmark data (in this case from the Australian Stock Exchange). The difference is then used as the probability for those eligible to transition this period. Those eligible would be non-shareowners of that sex, age and income.

A second method for alignment of probability transitions is to use the probability to form an ordered list. The probability of transition is calculated for individuals in scope based on their current circumstances. This probability is then used to rank the individuals. The number that makes the transition is set externally and this number is selected from the probability-ordered list of candidates. This is the method described on page 10 and used to align labour force movements.

Alignment of survival functions is done in one of three ways:

- Adjust the function using trial and error. A simple but laborious method of aligning the simulation outcomes is to add an adjustment to the baseline value in the equation. Estimating the value of the adjustment is not straightforward as the required results often vary over time and are exponential functions of the input values. This complexity limits the use of this technique. In DYNAMOD-3, the only one use of this method is made and that is in the alignment of first marriages.
- Adjust the function using a period-by-period calculation. This is similar to the trial and error method in that an adjustment is made to the baseline value of the survival function, except this time the adjustment is calculated dynamically. At the start of each period, a test run is undertaken using the unconstrained survival function. The results from this test run are compared with exogenous targets and an adjustment for that period is calculated. The actual selection is then done using the adjusted survival function.
- Use the survival function to rank rather than select. The survival function and random numbers are used to provide a ranking then a number are selected from this ordered list to achieve the desired outcome. The survival function is used as a ranking instrument rather than a selection mechanism. This technique is akin to the alignment of labour force transitions i.e. an exogenous target defines the size of the transition list and the function defines who is on that list. This technique is used to model fertility in DYNAMOD-3.

The combination of the move to a PC platform, the additional modules and the extended use of alignment is a model that is beginning to realise the potential offered by this dynamic microsimulation. The usefulness requirements of users – reliability, flexibility, and transparency – are being realised. The alignment of the simulation outcomes to external benchmarks and user-defined scenarios are making the biggest contribution to this progress. The biggest part within the alignment is the ability to control survival functions. It is this ability that will ensure DYNAMOD is used in the real world.

4.2 Exactly aligned fertility in DYNAMOD-3

With DYNAMOD-3, the method of aligning fertility has been changed from the approximate approach based on adjustment of the survival function to an exact alignment method where the survival function is used as a ranking mechanism. The survival function is used to generate a list of potential births, ranked in order of likelihood, with the birth events then drawn from this list until the target number have been selected. The advantages of this method are seen to be:

- the outcomes are aligned exactly (selection is made from the ranked list until the required number of events has been obtained);
- the method is flexible (different scenarios can readily be constructed by altering the targets);
- the methodology is simple (and accordingly more transparent to potential users);
- alignment is now being done against more readily available cross-sectional data rather than the longitudinal data normally required for survival functions; and
- the selection of people to experience the event is still influenced by our understanding of behaviour (the rank list is ordered with reference to how soon the person is predicted by the survival function to experience the event – those predicted to experience the event soonest are the first selected).

A potential constraint is the greater computer resources required to build and sort lists and run a selection process than was the case with the approximate alignment used in DYNAMOD-2. The resource demands are much greater, although continuous advances in processing power minimise the impact.

An important implication of this alignment method is that individual behaviour, while still important, is now taking a secondary role to alignment. In DYNAMOD-1, a person with characteristics that predicted a birth event to occur some years into the future would definitely not give birth in the next 12 months (unless they experienced some change in the background variables). Under the approximate alignment of fertility in DYNAMOD-2, it was unlikely that the adjustment to the survival function would bring the event forward more than a year or two. In DYNAMOD-3, however, if there is a requirement for a relatively large number of births to occur in a given period, then a person who was predicted to not give birth for many years could conceivably be selected to give birth in the period. Alignment now plays a much greater role in determining the modelled fertility outcomes.

That is the general picture and we now turn to a more detailed account of the alignment of fertility in DYNAMOD-3. We have seen that a number of limitations remained with the fertility alignment in DYNAMOD-2. These included:

1. the problem stemming from possible differences in the age distributions of the model population and actual population;
2. the point that different scenarios could not always easily be accommodated; and
3. the bias toward pre-marital births.

These issues have been addressed with DYNAMOD-3. To address the first two – the distribution of mother's age and scenario-setting issues – fertility rate targets were substituted for the numerical birth targets. In the previous version, for each year of the simulation, a target number of births for each of the seven age groups were read from a data file. In DYNAMOD-3, rather than a target, a fertility rate (births per 1000 females) for each age group is read from the data file. Clearly, if the numbers of females in each age group were the

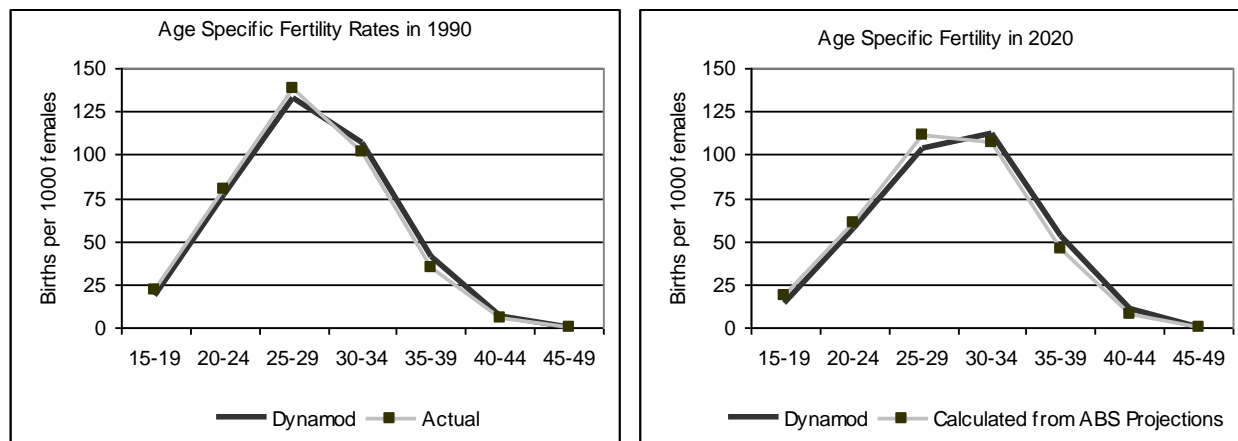
same in both the model and actual (target) populations, then this change would have no effect. However, the power of this change is in its ability to maintain age-specific birth rates with a changing age structure. The ease of specifying scenarios – which, in fertility terms, are typically thought of as age-specific fertility rates – has also increased dramatically.

The third issue – selecting too heavily from one type of birth – was addressed by using the ranking selection algorithm described above. Under the previous version those in scope for one of the three birth types was selected and a date of birth was determined from the adjusted survival function duration. If the date of birth was found to be in the forthcoming 12 months it was programmed to occur by entering the event into the crystal ball. If the date of birth was outside this period, it was discarded. In practice, females in the pre-marital group had a higher probability of pregnancy and childbirth. DYNAMOD-3 uses the survival function to **rank** those in scope for pregnancy but the function is not used to select individuals. Based on birth type proportions and age specific fertility rates read from the data file (which is normally based on ABS data but can be changed to model different scenarios), target birth numbers are calculated for the 21 sub-groups (7 age-groups * 3 birth types) on an annual basis. Females falling in scope for birth are divided into one of the 21 sub-groups and the survival function for that individual is calculated and recorded. The survival function produces a time to pregnancy based on the characteristics of the person and a random number. The survival function ‘score’ is used to rank that individual on her type and age-specific list. The target numbers of pregnancies for that sub-group are then selected from the ranked lists.

Another issue with age group alignment was identified. The selection of which month the birth would take place was initially selected at random. However, this resulted in a significant number of the mothers experiencing a change in their age group before the child was born. This impacted on the simulated age specific birth rates, particularly for the 15-19 age group. In this age group, the premarital birth survival function favours pregnancy for older women, i.e. there is a much higher chance of a 19-year-old woman becoming pregnant than a 15-year-old. The function produces the outcome that a high proportion of women selected to become pregnant in the next 12 months are aged 19. The random birth month approach resulted in a large number of these mothers turning 20 before giving birth and therefore being counted in the older age group. With random births the 15-19 age group birth rate could not be aligned. To overcome this problem, births were programmed to occur in the month before the mother’s birthday. Selection of this month ensured that the mother was still the same age at childbirth as when she was selected. The distribution of births is uniformly distributed over the year (assuming the mother’s birthdays are randomly distributed) but the mother-child distribution is constant. However, this is not believed to be of any real concern.

The aggregate result for births under DYNAMOD-3 maintained the high level of alignment achieved with previous version (see Figure 2 for the level of alignment). The improvements only become evident at the lower disaggregated age-specific fertility rate level. Whereas the previous version had not performed well at this level, DYNAMOD-3 showed good alignment between the simulated and benchmark data (Figure 3).

Figure 3: Age Specific Fertility Rates 1990 and 2020



Source: ABS Births Australia 1998 Cat no. 3301.0

In summary, the major changes that have been incorporated in the DYNAMOD-3 fertility alignment are the use of fertility rates in place of birth numbers; a change to the algorithm for calculating the month of childbirth; and a major change to the way in which survival functions are used. The survival functions are now used as a sorting tool rather than a selection method. This method increases the alignment of the model while the use of fertility rates rather than birth targets increases the flexibility.

5. Overall model performance

In the introduction, it was noted that a model that is not perceived as credible by its users, is of very limited value. To become useful, a model needs to be developed in such a way that it is plausible, flexible, transparent, and meets users' needs. The changes that have been implemented in DYNAMOD-2 and DYNAMOD-3 have been designed to fulfil these requirements. By aligning the outcome data with known reference data, the model is more plausible. The alignment mechanisms similarly provide the capacity for scenario setting and increases flexibility. Other steps, discussed below, have been taken to improve transparency.

Dynamic microsimulation models are complex. One reason for the complexity is the range and interaction of influences on an event. This interaction, while adding to the complexity, is one of the great strengths of dynamic modelling. For example, consider the changing of labour force status (employed, unemployed or not in the labour force). A transition is influenced by age, marital status, industry, occupation, recent earnings, season of the year and macro-economic data. In addition labour force status has an influence on childbirth, leaving home, couple formation, marriage, couple dissolution, returning to education, overseas emigration, holding two jobs, retirement pension, and annual earnings. A change to labour force status will influence all of these other characteristics of a person to varying degrees at different times.

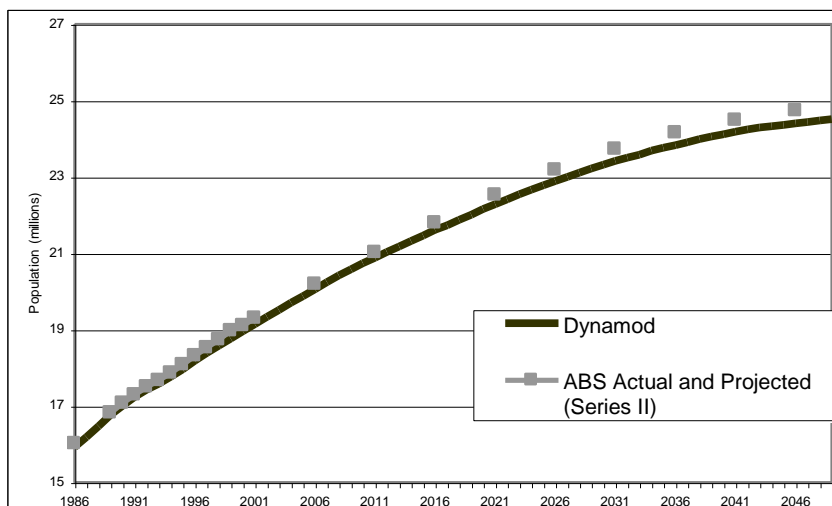
This interaction of influences makes transparency for the casual user or observer very difficult to achieve. However, a general understanding of how the model works and the interaction of the influences on behaviour can be provided in a relatively straightforward manner. In recent years, NATSEM has put substantial effort into documenting the processes within DYNAMOD using language that is user friendly (Kelly, 2000b). The documentation typically describes in a few paragraphs how the model works and then provides a one-paragraph

explanation of the influences on a particular transition. Technical language is avoided where possible.

The complexity also makes verification of the model difficult. Verification is concerned with ensuring that the model is actually doing what was intended and therefore involves checking the model both from a syntax perspective and from a logic perspective. Again, considerable effort has gone into checking the inputs, the processing and the output of each module of the model.

Perhaps the clearest path to establishing the model's credentials is, however, by demonstrating the model outcomes. Are they plausible and realistic? As an example, Figure 4 compares the actual total Australian population for the period 1986-2000 and the most probable projection of the 2001-2050 population of Australia (ABS Series II) with simulated outcomes. At this level, the model looks good.

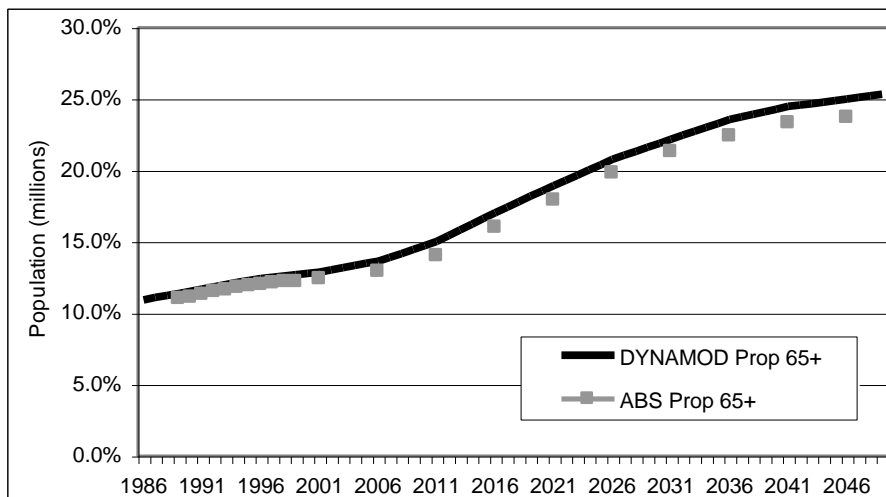
Figure 4: Australia's Population 1986 and 2050



Sources: ABS Population Projections, Cat no.3222.0

What if we look a little bit more closely at the outcomes – at the projected age structure of the population. The increasing size of the aged population and its consequences is a major policy issue in Australia, as in many other countries. In Figure 5 the rising population share of this group is presented. Again, actual and projected ABS data are plotted alongside DYNAMOD data and the result is a good match. Of course, the role of a dynamic microsimulation model is to generate more than these aggregate outcomes – which can be generated with far simpler tools. A demonstration of the performance needs to get down to showing how well (how plausibly and realistically) it generates individual lifetimes. The good performance at the broader level of analysis is, however, an important consideration and it is here that alignment has played a particular role in DYNAMOD.

Figure 5: Proportion of Australia's Population Aged 65+ 1986-2050



Sources: ABS Population Projections, Cat no.3222.0

6. Conclusion

Microsimulation development at NATSEM has proceeded on two fronts - static and dynamic. The static model has been operational for some years and has gained a reputation for providing useful and credible output. DYNAMOD, a dynamic model, has still not gained full credibility with its users. While the slower pace of acceptance for dynamic modelling is not totally unexpected, given the greater complexity, one area that has clearly hampered the model's acceptance is its capacity for alignment. In an attempt to address this and gain user acceptance, each new phase of development has included steps to strengthen alignment.

Development of the original version of DYNAMOD (1993-94) focused on the incorporation of techniques – survival functions – to represent individual behaviour more fully than was possible using the conventional transition probability approach. It was perhaps not appreciated at this time that while survival functions offered an elegant solution, they have disadvantages including inflexibility and being very difficult to align. In this early development environment, these disadvantages were not considered significant as the work focused on providing detailed individual data. Alignment of the output from the model with user-defined scenarios or with the real world was, at best, a secondary consideration.

Incorporation of alignment was a major focus of work in the second version of DYNAMOD (1995-98). The importance of alignment had been recognised and this phase of development saw the addition of alignment capabilities to much of the model, including some early attempts to align survival function processes. New modules which were added to the model in this version included alignment as basic elements.

The current version of DYNAMOD (1999-) has seen considerable increase in the sophistication and coverage of the alignment of the survival function processes. In some cases, the role of survival functions has changed from one of selection to ranking. Rather than use the output from the function to decide whether a transition is made, the function acts as a ranking tool and then a predetermined number are selected according to the rankings. This change increases the flexibility of the model and ensures alignment. DYNAMOD-3 also added more modules (which incorporate alignment) and better documentation.

The additional modules, the better documentation and the extensive use of alignment are significant steps and now allow DYNAMOD to claim its projections are useful. The outcomes have plausibility, flexibility and transparency and they align with either real or user-defined outcomes.

Bibliography

- Antcliff, S. 1993, *An Introduction to DYNAMOD*, DYNAMOD Technical Paper no. 1, National Centre for Social and Economic Modelling, University of Canberra.
- , Bracher, M., Gruskin, A., Hardin, A. and Kapuscinski, C. 1996, *Development of DYNAMOD: 1993 and 1994*, Dynamic Modelling Working Paper no. 1, National Centre for Social and Economic Modelling, University of Canberra.
- ABS (Australian Bureau of Statistics) 1998, *Population Projections 1997 to 2051*, ABS Catalogue No. 3222.0, Canberra.
- Bækgaard, H. 1996, ‘Calibration of earned income in DYNAMOD’, Research note, National Centre for Social and Economic Modelling, University of Canberra.
- 1997, ‘Dynamic microsimulation: the dynamic model structure revisited’, Note prepared for technical meeting on DYNASCAN, Available from National Centre for Social and Economic Modelling, University of Canberra, August.
- 2001, *Micro-Macro Linkage and the Alignment of Transition Processes: Some Issues, Techniques and Examples*, Technical Paper no. 25, National Centre for Social and Economic Modelling, University of Canberra, forthcoming.
- and King, A. 1995, ‘The relative merits of periodic transition probabilities and survival functions in dynamic microsimulation’, Research note, National Centre for Social and Economic Modelling, University of Canberra.
- Bracher, M. 1993, ‘Preliminary Demographic Parameters for the DYNAMOD micro-simulation model’, in Antcliff et al, *Development of DYNAMOD: 1993 and 1994*, Dynamic Modelling Working Paper no. 1, National Centre for Social and Economic Modelling, University of Canberra, pp 37-78.
- Galler, H. 1997, *Discrete-Time and Continuous-Time Approaches to Dynamic Microsimulation Reconsidered*, Technical Paper no. 13, National Centre for Social and Economic Modelling, University of Canberra
- Gruskin, A. 1994, ‘Dynamic Microsimulation: A programming approach called DYNAMOD’, in Antcliff et al, *Development of DYNAMOD: 1993 and 1994*, Dynamic Modelling Working Paper no. 1, National Centre for Social and Economic Modelling, University of Canberra, pp 1-36.
- Hellwig, O. 1995, ‘Issues in NATSEM’s development of a dynamic microsimulation model’, National Centre for Social and Economic Modelling, University of Canberra, November.
- Kelly, S. 2000a, ‘Modelling the impact of Housing Values on the Wealth of Australians’, Paper presented to the 26th General Conference of the International Association for Research in Income and Wealth, Poland, September.
- , 2000b, *Overview of DYNAMOD-3*, available from National Centre for Social and Economic Modelling, University of Canberra.

- , Percival, R. and Harding, A. 2001, ‘Women and Superannuation in the 21st Century: Poverty or Plenty?’, Paper presented to SPRC National Social Policy Conference 2001, University of NSW, July.
- King, A. 1995, ‘Dynamic Microsimulation modelling at NATSEM: Where are we and where to from here?’, National Centre for Social and Economic Modelling, University of Canberra.
- and Bækgaard, H. 1997, ‘DYNAMOD: The Experience and Lessons’, Note prepared for Uppsala meeting on DYNASCAN, National Centre for Social and Economic Modelling, University of Canberra.
- , Bækgaard, H. and Robinson, M. 2000, *DYNAMOD-2: An Overview*, Technical Paper no.19, National Centre for Social and Economic Modelling, University of Canberra.
- , Walker, A. and Bækgaard, H. 2000, *Modelling Overseas Migration in DYNAMOD-2*, Technical Paper no.22, National Centre for Social and Economic Modelling, University of Canberra (forthcoming).
- Morrison, R., 2000, ‘Assessing the Quality of DYNACAN’s Synthetically-Generated Earnings Histories’, Social Policy Directorate, Human Resources Development Canada, April.
- Robinson, M. and Bækgaard, H. 2001, *Modelling Students in DYNAMOD-2*, Technical Paper no.23, National Centre for Social and Economic Modelling, University of Canberra, forthcoming.
- Sargent, R.G. 2000, ‘Verification, Validation and Accreditation of Simulation Models’, in Jones et al (ed.), *Proceedings of the 2000 Winter Simulation Conference*, pp 50-59.