Can population projections be used for sensitivity tests on policy models?

John Bryant

NEW ZEALAND TREASURY

WORKING PAPER 03/07

JUNE 2003

THE TREASURY
Kaitohutohu Kaupapa Rawa

NZ TREASURY WORKING PAPER 03/07	Can population projections be used for sensitivity tests on policy models?
MONTH/YEAR	June 2003

AUTHOR

John Bryant
The Treasury
PO Box 3724
Wellington
New Zealand

Email	John.Bryant@treasury.govt.nz
Telephone	+64 4 917 7027
Fax	+64 4 473 1151

ACKNOWLEDGEMENTS

John Creedy, Dharma Dharmalingam, John Janssen, and Veronica Jacobsen provided helpful comments on earlier versions of this paper.

NZ TREASURY

New Zealand Treasury
PO Box 3724
Wellington 6008
NEW ZEALAND

Email	information@treasury.govt.nz
Telephone	64 4 472 2733
Website	www.treasury.govt.nz

Abstract

Many policy models require assumptions about future population trends. Sensitivity tests for these assumptions are normally carried out by comparing population projection variants. This paper outlines some of the conditions that variant-based sensitivity tests must meet if they are to be informative. It then describes four common situations where these conditions are not met, so that conventional sensitivity tests are not informative. The solution, the paper argues, is stochastic population projections.

JEL CLASSIFICATION C520 - Model Evaluation and Testing

E170 - Forecasting and Simulation

J110 - Demographic Trends and Forecasts

KEYWORDS Demography; Sensitivity testing; Population projections; Policy modelling

Table of Contents

Abstract i

Table of Contents ii

List of Tables ii

List of Figures ii

1 Introduction 1

2 A framework for assessing demographic sensitivity tests on policy models 2

3 Fertility, mortality, and migration assumptions exclude random shocks 5

4 'Low-low or 'high-high' variants are not calculated 6

5 Variants' rankings differ with the outcome considered 7

6 Fertility only changes early in the projection period 8

7 A solution: Stochastic population projections 10

8 References 12

List of Tables

Table 1 - Coverage of possible scenarios and the implications for sensitivity testing 3
Table 2 - Interpretation of results from a sensitivity test 4
Table 3 - Statistics New Zealand projections variants for total population and dependency ratios in 2021 8

List of Figures

Figure 1 - Demographic sensitivity tests for policy models 2
Figure 2 - Estimates and 1999-base projection assumptions for net permanent and long term migration into New Zealand 5
Figure 3 - Simulation of the effects of a fertility decline 8
Figure 4 - Statistics New Zealand fertility assumptions and projected values for percentage of population aged 15-64, 1999(base)-2101 projections 9

Can population projections be used for sensitivity tests on policy models?

1 Introduction

Many policy models require projections of future population size and structure. Macroeconomic models that include variables for the labour force or the size of the tax base, for instance, require data on the size and age-distribution of the working-age population. Forecasts of future needs for hospitals, schools, or prisons all require data on potential occupants.

Sometimes modellers use only one projection variant, typically the ‘central’, ‘median’, or ‘medium’ series prepared by the relevant statistical agency. Often, however, modellers require some indication of how uncertainty concerning the demographic variables affects the robustness of the model results. The standard tool for doing so is projection variants.

Projection variants are generated by varying assumptions about future paths for fertility, mortality, and migration. The status and interpretation of the projection variants is often ambiguous. Demographers are generally unwilling to attach explicit probabilities to the variants, and commentaries on the projections often warn the reader that the variants are hypothetical scenarios rather than predictions. However, the commentaries often refer to some variants as more plausible than others, or state that certain events, such as the population reaching a given level, are ‘likely’.

Most modellers appear to take a pragmatic stance towards these conceptual ambiguities. They enter the variants into their policy models, and compare the outcomes. If the outcomes are similar for all variants, modellers state that their forecasts are insensitive to demographic uncertainty. If the outcomes differ, modellers warn their readers accordingly and call for further research. Few modellers give any indication that they are dissatisfied with this situation.

This paper argues that the conventional approach is seriously flawed. It presents examples in which population variants provide a misleading indication of uncertainty about demographic variables. The paper explores the underlying reason for these problems, and argues that they prevent effective sensitivity testing.

The paper considers only national projections. Projections for groups of countries or for regions within countries involve addition difficulties, discussed in Lee (1998: 164-5), Bongaarts and Bulatao (2000: 198), and Siegel (2002: 460-82).

Section 2 of the paper sets out a framework for assessing the value of demographic sensitivity tests with policy models. The four subsequent sections describe specific problems. Section 3 describes how the omission of random shocks from trajectories for fertility, mortality, and migration leads to actual population sizes exceeding all projected population sizes soon after the projections are published. Section 4 describes how the absence of low-fertility, low-mortality variants and high-fertility, high-mortality variants reduces the range covered by projected dependency ratios. Section 5 looks at how variants that bracket a substantial range for one population variable may bracket only a narrow range for another population variable. Section 6 examines how the practice of restricting fertility changes to the beginning of the projection interval can lead to confusing results for trends in age structure. The paper concludes with a discussion of stochastic population projections, which are a promising alternative to the variants approach.

2 A framework for assessing demographic sensitivity tests on policy models

Figure 1 shows the steps involved in demographic sensitivity testing. The population projections are generally carried out by the relevant statistical agency. Future paths for fertility, mortality, and migration are chosen. These are entered into a population projection model, such as the standard 'cohort components' model,[1] and future paths for population size and structure are derived. These paths are the 'population variants' referred to in the population projections literature. Complete descriptions of the variants consist of variables giving the size of each age-sex group, in each year of the projection. Many derived variables are, however, produced, such as dependency rates, numbers of school-age children, or total population size.

Figure 1 - Demographic sensitivity tests for policy models

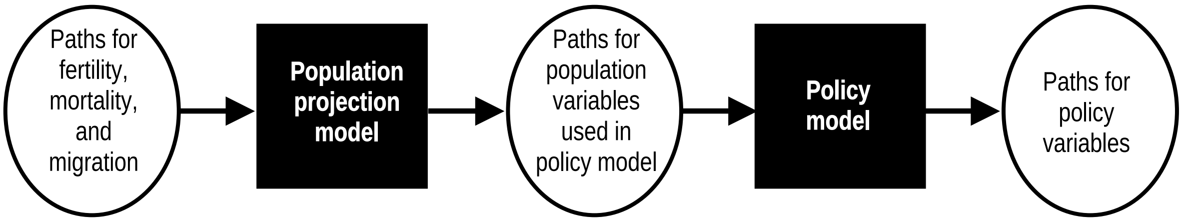

Users of policy models generally take the population variants as given. Some models require all the detail produced in the population projections. Typical health expenditure models, for instance, require population numbers for every age-sex group. Other models require only a few derived variables. Some macroeconomic models, for instance, require nothing more than numbers for the total and working-age population. Paths for the required variables are entered into the policy model and the results compared, in an attempt to learn something about the sensitivity of the model's results to demographic uncertainty.

[1] Typically, a path for fertility, mortality, or migration is specified using a single variable: mortality paths, for instance, are often specified using life expectancy. To carry out population projections, entire schedules of age-sex-specific rates are needed. These schedules are derived from a model relating overall levels to underlying rates (see, for instance, Lee and Carter 1992). The discussion in this paper implicitly treats such models as part of the overall population projection model.

Under what conditions does this procedure in fact provide informative results? Table 1 presents a simple typology of cases arising during sensitivity testing, and shows admissible conclusions under each case.

Table 1 - Coverage of possible scenarios and the implications for sensitivity testing

	Coverage of empirically possible scenarios			Conclusion about sensitivity of policy model to demographic assumptions
Case	Fertility, mortality, and migration	Population variables used in policy model	Outcome variables from policy model	
1	Wide	Wide	Wide	Sensitive
2	Wide	Wide	Narrow	Insensitive
3	Narrow	Wide	Wide	Sensitive
4	Narrow	Wide	Narrow	Insensitive
5	Narrow	Narrow	Wide	Sensitive
6	Narrow	Narrow	Narrow	No conclusion possible

In Case 1, coverage of empirically possible scenarios for fertility, mortality, and migration is wide. In other words, the sets of fertility, mortality, and migration assumptions that are entered into the population projection model in Case 1 jointly cover a broad range of plausible conditions. In Case 1, coverage of possible scenarios for population variables is also wide. This is likely when coverage of possible fertility, mortality, and migration scenarios is also wide, since changes in population size and age structure are completely determined by changes in fertility, mortality, and migration. Finally, in Case 1, coverage of possible outcomes from the policy model is wide. Entering different population scenarios into the policy model gives substantially different outputs. The correct conclusion is that the policy model's results are sensitive to demographic assumptions. Uncertainty over future demographic variables carries through to uncertainty about the model results.

Case 2 is identical to Case 1, except that the range of outcomes from the policy model is narrow. Entering different population inputs into the policy model has little effect on the model outcomes, even though the population inputs cover a wide range of possible cases. The model user is entitled to infer that the model is insensitive to demographic assumptions. This is the result modellers generally prefer.

In Cases 3 and 4, coverage of possible fertility, mortality, and migration scenarios is narrow. Coverage of possible scenarios for the population variables used in the policy model is, however, wide. This combination of wide and narrow coverage does arise in practice. One example is when low and high migration assumptions differ markedly, and the only population variable used in the policy model is total population size. Wide and narrow coverage of the outcomes from the policy model lead to the same conclusions about the sensitivity of the model in theses cases as they do in Cases 1 and 2.

In Case 5, the narrow coverage of fertility, mortality, and migration carries through to coverage of population variables. Coverage of policy model outcomes is, nevertheless, wide. The fact that outcomes from the policy model vary substantially even when the population scenarios vary relatively little implies that the model is definitely sensitive to demographic uncertainty.

Finally, in Case 6, coverage is narrow for fertility, mortality, and migration, and for the population variables, and for model outcomes. The narrow coverage of model outcomes is consistent with the model being insensitive to the demographic assumptions, but it may simply reflect the fact that the population scenarios entered into the model covered only a small proportion of the plausible range. No conclusion about the policy model's sensitivity to demographic assumptions is therefore possible.

Cases 1 and 2, in which wide coverage of possible fertility, mortality, and migration scenarios ensures wide coverage of possible population scenarios, never occurs when working with population variants. As later sections of this paper illustrate, this is because the set of plausible scenarios for fertility, mortality, and mortality trajectories is too large and multi-dimensional to be adequately represented by a small number of population variants.

Work with population variants only leads, then, to Cases 3-6. Whether or not modellers draw the correct conclusions from these cases depends on which cases the modellers believe to have occurred. The extent to which the policy model produces a wide range of outcomes during demographic sensitivity testing is readily observable. Modellers who encounter Cases 3 or 5 and observe a wide range of outcomes are therefore likely to assume that one or other of these cases has occurred; if the modellers are unfamiliar with the limitations of population variants, they might also assume that Case 1 has occurred. Similarly, modellers who encounter Cases 2 or 4 and observe a narrow coverage of possible outcomes are likely to assume that one of Cases 2, 4, or 6 has occurred.

Table 2 - Interpretation of results from a sensitivity test

Case that actually occurred	Case that modeller believes occurred	Modeller's interpretation of the sensitivity test
3 or 5	1, 3, or 5	Correctly concludes that model sensitive to demographic uncertainty
4	2 or 4	Correctly concludes that model insensitive to demographic uncertainty
4	6	Incorrectly concludes that test uninformative
6	2 or 4	Incorrectly concludes that model insensitive to demographic uncertainty
6	6	Correctly concludes that test uninformative

Table 2 shows the possible combinations of cases and modellers' beliefs, and the consequences for the correctness of the modellers' interpretations. The first row of the table shows what happens when Cases 3 or 5 occur. Regardless of whether the modellers believe that Case 1, 3, or 5 has occurred, they still conclude, correctly, that the policy model is sensitive to demographic uncertainty.

The second and third rows show combinations occurring under Case 4, when the policy model produces only a narrow range of outcomes. If the modellers assume that Cases 2 or 4 have occurred, they conclude, correctly, that the model is insensitive to demographic uncertainty. If, however, the modellers assume that Case 6 has occurred, so that the narrow range of outcomes is simply a result of a narrow range of population scenarios, they conclude, incorrectly, that the test is uninformative.

Finally, the fourth and fifth rows of the table show combinations occurring under Case 6, when the range of population scenarios and model outcomes are both narrow. If modellers overlook the narrow range of population outcomes, and assume that Cases 2 or 4 have occurred, they conclude, without warrant, that the model insensitive to demographic uncertainty. If they assume that Case 6 has occurred they reach the correct but unhelpful conclusion that the test is uninformative.

Case 6 is evidently the most problematic. It does, however, often arise in practice. The following four sections describe some examples.

3 Fertility, mortality, and migration assumptions exclude random shocks

Time series for fertility, morality, and migration can in principle be separated into trend and random shock components. Figure 2, for example, shows the time series for net annual permanent and long term (PLT) migration to New Zealand between 1970 and 2002.[2] The trend seems to be somewhere around zero, with large random shocks around this value.

Figure 2 - Estimates and 1999-base projection assumptions for net permanent and long term migration into New Zealand

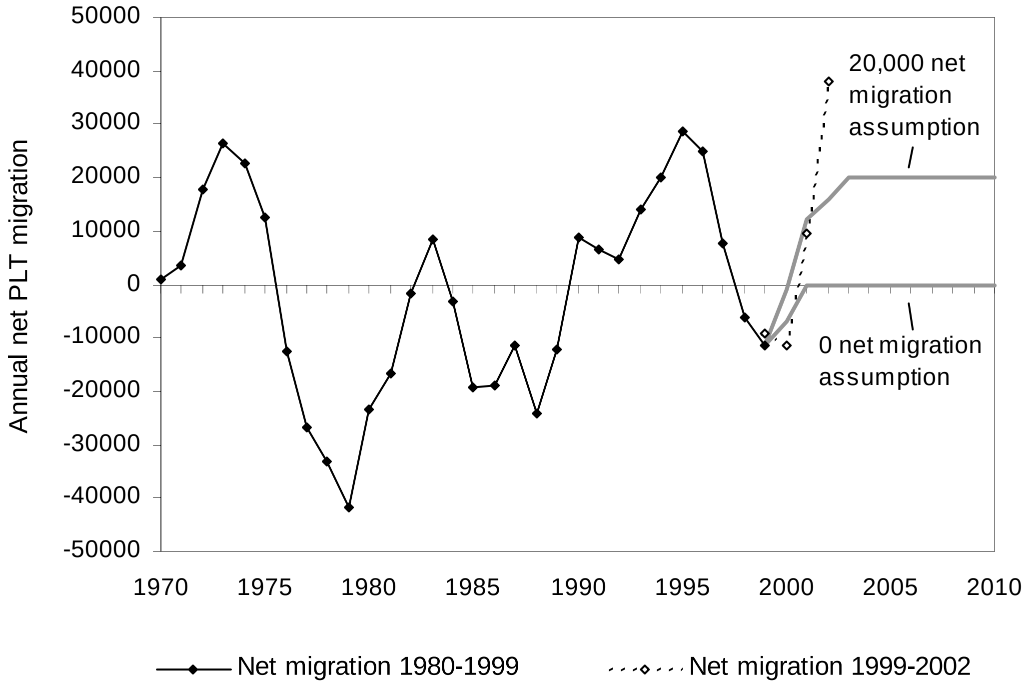

Source – All data obtained from documents on the Statistics New Zealand website www.stats.govt.nz. Estimates for migration 1970-1979 obtained from 'Demographic Trends 2001', Table 5.01. Estimates for migration 1980-1998 obtained from 'Demographic Trends 2001', Table 5.01. Estimate for 1999, and description of migration assumptions obtained from 'National Population Projections, 1999(base)-2101', Table 3.01. Estimates for 1999-2002 obtained from 'Key demographic indicators, 1999-2002'. The 1999 estimates from this and the previous source differ slightly.

The most natural interpretation of the assumptions used in the construction of population projections is that random shocks have been excluded (Lee 1998: 156; Bongaarts and Bulatao 2000: 191). Figure 2 again provides an example. The figure shows the first 11 years for two migration assumptions used by Statistics New Zealand to construct

[2] Permanent or long term migrants, as distinguished from short-term visitors, are people entering New Zealand who state on their arrival cards that they intend to remain in the country for at least 12 months, or people leaving New Zealand who state on their departure cards that they intend to remain outside the country for at least 12 months.

projections for the period 1999-2101.[3] The '0' assumption is Statistics New Zealand's lowest migration assumption, and the '20,000' assumption its highest. Values over the first few years of each assumption reflect the fact that the migration in the launch year differed from assumed trend level. In subsequent years, however, the values are fixed at trend levels. There is none of the volatility apparent in the historical series. Using the typology introduced in the previous section, coverage of possible paths for fertility, mortality, and migration is 'narrow.'

Virtually all statistical agencies omit random shocks from series for mortality, fertility, or migration. Omission of random shocks is essentially unavoidable when using the population variants approach. If only a handful of projection variants are calculated, and if shocks can take an indefinitely large number of forms, then having no shocks is the least arbitrary choice.

Omitting random shocks from population projections does not greatly reduce their coverage of plausible figures for long-run population size. It does, however, have clear and rather awkward consequences for short-run projections. As Figure 2 illustrates for migration, high and low assumption about fertility, mortality, and migration typically differ little in the early years of a projection. Moreover, because little time has passed, the effects of any differences have yet to accumulate. High and low variant projections for outcomes such as population size therefore cover a narrow range.

Random shocks to fertility, mortality, or migration early in a projection period often cause population size and other variables to fall outside this range. As Figure 2 shows, net migration to New Zealand experienced a large positive shock over the three years following the 1999-base projections. By 2002, migration levels were well above the trend level implied by the '20,000' assumption. This is one reason why the New Zealand population (probably) reached 4 million by April 2003, although none of the projection variants calculated in 1999 had the New Zealand population attaining 4 million until early 2004. (Another reason for the early arrival at 4 million was that estimates of population size in the base year for the projection turned out to be too low.) This situation is certainly not unique. Around the world, actual outcomes frequently fall outside the range set by high and low projection variants soon after the projections are published (Lee 1998: 156).

Policy models that use short-term changes in population size as an input include funding formulas for health and education. Health and education budgets in New Zealand have, indeed, been revised in the wake of the recent, unexpectedly large, population increases. The revisions have been greater than would have appeared likely from demographic sensitivity testing at the time when the 1999-base projections were prepared. Demographic sensitivity tests at this time would have fallen under Case 6 of the typology set out in Table 1.

4 'Low-low or 'high-high' variants are not calculated

Even when time series for fertility, mortality, and migration are each restricted to 3-4 different alternatives, a statistical agency that wanted to cover all possible combinations of these series would need to produce 27-64 projection variants. In practice, statistical

---

[3] This paper uses the 1999-base projections rather than the more recent 2001-base projections in this and subsequent examples because, at the time of writing (May 2003), detailed series over a 100-year period for the 2001 base are not yet publicly available, and because the use of 1999-base projections permits comparison between projected and actual results.

agencies produce only a fraction of the possible number. Statistics New Zealand, for instance, produces eight different projection variants, and the United Nations Population Division produces four.

Many of the possible variants that are not produced are nevertheless plausible. Neither Statistics New Zealand nor the UN Population Division routinely publish variants combining their low fertility and low mortality assumptions. Yet it is possible that fertility and mortality rates will both be lower in future than is currently expected: in developed countries, fertility and mortality are at present much lower than most demographers were predicting 30 year ago.

Some of the variants that are not produced but are nevertheless plausible give more extreme results than the variants that are produced. The combination of low fertility and low mortality, for instance, yields a higher ratio of old people to young people than any other variant. The combination of high fertility and high mortality, which is also not generally calculated, yields a lower ratio of old to young than any other variant. Omitting low-low and high-high variants leads to an artificially narrow range for ratios between old and young.

Forecasts of health expenditures per capita are significantly affected by the ratio of old to young, since old people attract much greater expenditure per capita than young people. Forecasts of pension expenditures are, of course, even more strongly affected. The potential importance of demographic trends is, however, understated when the range for the ratios between old and young is narrow. This is another instance of Case 6.

5 Variants' rankings differ with the outcome considered

Changes in fertility, mortality, and migration rates all have different effects on population structure. An illustration is provided by Table 3, which gives selected results from Statistics New Zealand's 2001-base population projections for 2021. As comparison of Series 2 and 6 shows, differences in migration rates have a major effect on population size but have a minor effect on dependency rates. Comparison of Series 1 and 8 shows that, in contrast, differences in fertility rates have a minor effect on population size, but a major effect on dependency rates. An alternative choice of time period, or differences in the age-profile of migration, could affect the comparison. It is generally true, however, that variants do not have stable rankings, in that a variant yielding high values for one variable, such as population size, does not necessarily yield high values for another variable, such as the dependency ratio (Lee 1998; Bongaarts and Bulatao 2000: 192-4).

Modellers who are not aware of this phenomenon are at risk of constructing demographic sensitivity tests that are even weaker than necessary. Modellers might, for instance, be impressed by the ability of Series 1 and 5 in Table 3 to bracket a wide range of plausible dependency ratios and not realize that these series do not bracket a wide range of plausible population sizes. If the modellers use Series 1 and 5 to carry out a sensitivity test on a policy model that is sensitive to population size, this sensitivity may not be apparent. This is again an instance of Case 6.

Table 3 - Statistics New Zealand projections variants for total population and dependency ratios in 2021

Series	Fertility	Mortality	Net migration	Population, thousands	(As % of Series 4)	Dependency ratio	(As % of Series 4)
1	Low	Medium	5,000	4,396	(98%)	51%	(94%)
2	Medium	Medium	0	4,374	(97%)	55%	(102%)
4	Medium	Medium	5,000	4,506	(100%)	54%	(100%)
6	Medium	Medium	20,000	4,821	(107%)	53%	(98%)
8	High	Medium	5,000	4,616	(102%)	57%	(106%)

Source - Tables 1 and 3 of the Statistics New Zealand's *National Population Projections (2001(base) - 2051*) available online at www.stats.govt.nz

6 Fertility only changes early in the projection period

Figure 3 shows results from a simulation of the effects of fertility decline on the proportion of the population in the working ages. The fertility indicator used is the 'total fertility rate' (TFR), which is defined as the number of children the average woman would have over her lifetime if prevailing age-specific fertility rates were to be maintained indefinitely.[4] The TFR is constant at 2.15 for the first 20 years of the simulation, it declines to 1.65 over the next 20 years, and then again remains constant. Life expectancy (not shown) stays at 80 years throughout, and there is no migration.

Figure 3 - Simulation of the effects of a fertility decline

As can be seen in Figure 3, the effects of the fertility decline are complex. It initially drives the proportion in the working ages higher than it would otherwise have been, and then drives it lower. The reason for the rise and fall is essentially that a reduction in birth rates

[4] Equivalently, the total fertility rate is the sum of the single-year age-specific fertility rates.

reduces the growth rate of younger age groups before it reduces the growth rate of older age groups (Preston, Heuveline, and Guillot 2000: 165-7). Fertility increases have the opposite effect: they drive the proportion in the working ages lower than they would otherwise have been, before driving them back upwards.

Changes in fertility rates have strong effects like these on any ‘intermediate’ age groups—age groups that have some proportion of the total population that is younger than them and some proportion that is older. Changes in mortality rates can generate similar effects for intermediate age groups when the changes are concentrated in the the youngest age groups, but this no longer the case outside the least developed countries.

Figure 4 - Statistics New Zealand fertility assumptions and projected values for percentage of population aged 15-64, 1999(base)-2101 projections

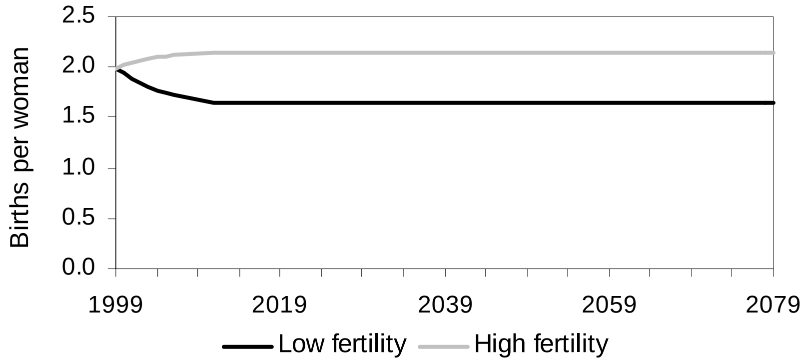

(i) Fertility assumptions

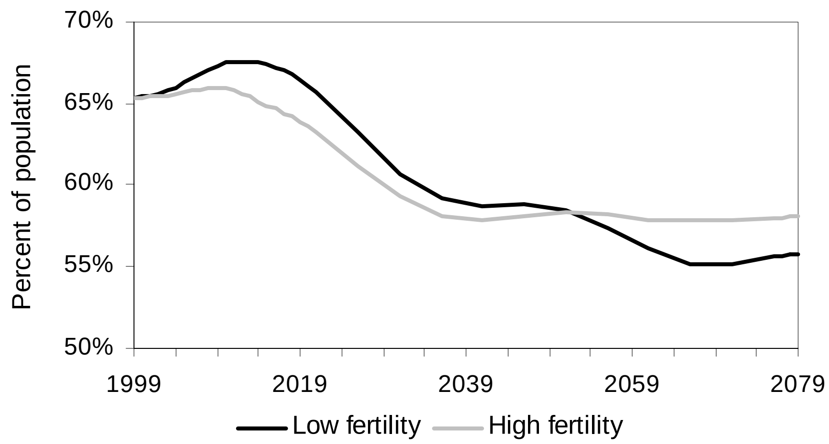

(ii) Percent of population aged 15-64

Source - ‘National Population Projections, 1999(base)-2101’, Tables 3.01, 4.01. and 4.08, downloaded in August 2002 from Statistics New Zealand website www.stats.govt.nz.

Fertility assumptions are typically constructed so that all the changes occur early in the projection period, as fertility moves from its initial level to its trend level. An example is given in the upper panel of Figure 4, which shows the fertility assumptions from Statistics New Zealand’s 1999-base projections. Confining change to the beginning of the projection period is probably the least arbitrary approach. It does, however, cause problems for sensitivity testing.

As the simulation results suggest, the distinctive age structures created by changes in fertility rates, with unusually large or unusually small concentrations of population in the intermediate age groups, last for only a generation or two. Confining fertility changes to the early part of the projection also confines these peculiar age structures to the early part of the projection period, even though, in reality, they can occur at any time.

A related problem is that projection variants fail, in a rather striking way, to cover the plausible range for the percentage of the population in any intermediate age group. Confining fertility changes to early in the projection period can lead to situations where projection variants based on high and low fertility levels converge over time, instead of diverging. With certain choices for mortality and fertility, and a sufficiently long projection period, the variants may even cross. The lower panel of Figure 4 provides an example. The low fertility variant and high fertility variant initially diverge, as the fall in fertility rates drives the proportion in the working ages up and the rise in fertility rates drives the proportion in the working ages down. By the 2020s, however, these effects begin to dissipate, and the variants start to converge. In the case depicted in Figure 4, the variants eventually cross. The crossover itself does not create difficulties for sensitivity testing; what matters is the narrow range between variants during the years before and after the crossover.

These problems can undermine attempts to conduct sensitivity tests on policy models that depend on assumptions about the size of intermediate age groups. Many models do in fact depend on such assumptions. The size of the tax base, for instance, depends on the proportion of the population in the working ages, and the number of potential army recruits depends on the proportion in the prime combat ages.

These problems do not, however, appear to be widely recognised. Sensitivity tests are often carried out with fertility variants that are likely to have entered their convergent phase. One representative example is provided by a widely-cited OECD (1998: 123) report on population ageing. In a text box entitled 'Different assumptions about demography make little difference', a bar chart demonstrates that, 35 years into the projection, the proportion of the population aged 15-64 is much the same in the low and high fertility variants. This is another instance of Case 6. The restricted range for the fertility assumptions, leads to a restricted range for population variables, which removes the value of the sensitivity test.

7 A solution: Stochastic population projections

In the typology set out in Table 1, Cases 3, 4, 5 can all be obtained using a variants-based approach to sensitivity testing, and all produce useful results. However, designing a sensitivity test so that these cases occur, and are known to occur, can be difficult. The previous four sections illustrate these difficulties.

There is, however, a promising alternative to variants-based testing. Over recent years, demographers have made considerable progress in developing stochastic population projections (Lee 1998; Lutz, Sanderson, and Scherbov 2001). Stochastic population projections all use some method for randomly generating large sets of realistic paths for fertility, mortality, and migration. Some methods apply time series methods to obtain means and variances (Lee and Tuljapurkar 1994), while others rely more on expert judgement (Lutz et al 2001). Results depend crucially on the covariance between different variables. The standard assumption is that age-specific rates for same variable, such as fertility, are perfectly correlated, while adjacent years are partly correlated, and distinct variables, such as fertility and mortality, are uncorrelated (Lee 1998).

The fertility, mortality, and migration trajectories are entered into standard population projection models, to produce large sets of population projections. Demographers summarize these sets by calculating means, variances, and confidence intervals for key variables such as population size and the dependency rate. Carrying out sensitivity tests on a policy model is simplest when the model requires only the key variables. In this case, users can simply enter values that, on the basis of the variance and confidence intervals, appear suitably extreme. Testing is more difficult with models that require highly detailed demographic inputs, such as models of health expenditure. Users may, in this case, need to enter the full set of population projections, rather than summary statistics.

Stochastic population projections can, accordingly, be unwieldy. They are also technically demanding, and, in the case of time series methods, require long series of historical data. Furthermore, existing methods for randomly generating fertility, mortality, and migration paths are still not entirely satisfactory. Even with time series methods, for instance, users still need to specify a long-term trend level for fertility (Lee and Tuljapurkar 1994). Some demographers argue, in addition, that the assumption of perfect correlations between age-specific rates can and should be relaxed (Booth, Maindonald, and Smith 2002).

Stochastic population projections do, however, allow the user to obtain Cases 1 and 2 of the typology in Table 1. These are the cases in which the plausible ranges for fertility, mortality, and migration, and hence the ranges for the population variables, are adequately covered. Adequate coverage of these ranges means that the results from sensitivity tests can be interpreted easily and safely. Stochastic population projections can put sensitivity testing on a surer footing.

Applications of stochastic population projections to important policy questions have begun to appear. The United States Congressional Budget Office (2001), for instance, has used stochastic population projections to forecast social security expenditures. The New Zealand Treasury has carried out similar work for social expenditures by the New Zealand government (Creedy and Scobie 2002). Variant-based population projections are still, however, more commonly used than stochastic projections.

Demographers sometimes try to promote greater use of stochastic projections by pointing out that stochastic projections have clearer conceptual status than variant-based projections, or by noting how stochastic projections can be incorporated into an elegant Bayesian decision-making framework (Tuljapurkar 1992). It seems unlikely, however, that practical minded users of policy models will be persuaded that these benefits outweigh stochastic projections' additional costs. Users of policy models may be more interested in the capacity of variant-based and stochastic projections to support meaningful sensitivity tests. On this measure, stochastic projections clearly outperform variant-based projections. This suggests that stochastic projections will become increasingly popular.

8 References

Bongaarts, J and R A Bulatao (2000) *Beyond Six Billion: Forecasting the World's Population*. (Washington DC: National Academy Press).

Booth, Heather, John Maindonald, and Len Smith (2002) "Applying Lee-Carter under conditions of variable mortality decline." *Population Studies* 56(3): 325-336.

Congressional Budget Office (2001) "Uncertainty in Social Security's Long-Term Finances: A Stochastic Analysis." Washington DC, Congressional Budget Office, Working Paper.

Lee, Ronald and Lawrence Carter (1992) "Modelling and Forecasting US Mortality." *Journal of the American Statistical Association* 87(419): 659-671.

Lee, Ronald D (1998) "Probabilistic approaches to population forecasting." in Wolfgang Lutz, James W Vaupel and Dennis A Ahlburg (eds) *Frontiers of Population Forecasting* (New York: The Population Council).

Lee, R and S Tuljapurkar (1994) "Stochastic projections for the United States: Beyond high, medium, and low." *Journal of the Americal Statistical Association* 89(428): 1175-1189.

Lutz, Wolfgang and Sergei Scherbov (1998) "An expert-based framework for national population projections: The example of Austria." *European Journal of Population* 14(1): 1-17.

Lutz, Wolfgang, James W Vaupel and Dennis A Ahlburg (1998) *Frontiers of Population Forecasting*. (New York: The Population Council).

Lutz, Wolfgang, Warren Sanderson and Sergei Scherbov (2001) "The end of world population growth." *Nature* 412: 543-545.

OECD (1998) *Maintaining Prosperity in an Ageing Society*. (Paris and Washington, DC: Organisation for Economic Cooperation and Development).

Siegel, Jacob S (2002) *Applied Demography: Applications to Business, Government, Law, and Public Policy*. (San Diego: Academic Press).

Statistics New Zealand, ‘Demographic Trends 2001: All Tables’. Document downloaded from Statistics New Zealand website www.stats.govt.nz in August 2002.

Statistics New Zealand, ‘National Population Projections, 1999(base)-2101: All Tables’ Document downloaded from Statistics New Zealand website www.stats.govt.nz in August 2002.

Statistics New Zealand, ‘Key demographic indicators, 1999-2002’. Page on Statistics New Zealand website www.stats.govt.nz. Document downloaded in April 2003.

Tuljapurkar, Shripad (1992) "Stochastic population forecasts and their uses." *International Journal of Forecasting* 8(3): 385-91.