Extended abstract

A wide variety of decisions on infrastructure planning and services provision depend on expectations about future changes in the size, structure, and spatial distribution of the population (e.g. Auerbach and Lee 2001). Therefore, accurate and timely population projections, and/or forecasts, are an essential input into the decision-making of local and national government planners, private sector developers, businesses, community and advocacy groups, and others. Despite the critical nature of population projections for decision-making, there is no one universally adopted methodology, although the cohort-component projections model, that dates back to Whelpton (1928), has remained until the present the most popular technique – at least for relatively large populations. However, even with a given methodology, there can be many differences in model equations, parameters, and underlying assumptions. The choice of a population projections model and its specification depends on the goals of the modeller in terms of the desired validity and accuracy of the projections, but in practice also on the acceptability of the model and its projections to decision-makers and other end-users.

That the acceptability of model outputs by end-users, who are usually not modelling experts, is important for deciding on a projection methodology may strike some as problematic. However, if end-users cannot accept that the model equations, parameters and assumptions reflect the reality 'on the ground', then they are less likely to adopt the projections as an input into their decision-making, and the utility of the resulting projections will be reduced. Thus 'face validity', i.e. when the projections model and its results seem reasonable to the users, is an important test of any population projections exercise, although highly subjective (Sargent, 2013).

One threat to the face validity of a population projections exercise is the 'black box' nature of some projection models. If the model and equations are too complex for modellers to adequately explain, or for end-users to understand or interpret, then the acceptability of the model will be reduced. This problem is exacerbated in the case of emerging technologies such as machine learning algorithms, which have now started to be adopted in sub-national population projections (the recent survey by Wilson et al. 2021 provides several examples). Forecasts derived by machine learning are likely to be seen as truly 'black box' in nature by many end-users. Similarly, if the model equations do not appear to incorporate the drivers of population change that end-users expect that they will include, then the outputs of the model may be dismissed as not capturing important features of population change.

As noted already above, the most commonly employed method for projecting the population at both national and regional levels is the cohort-component model. This model has not only gained its popularity by the relative accuracy and usefulness of its outputs, but also by the way in which it makes population change transparent to end-users by explicitly quantifying the components of change (births, deaths and migration).

In terms of face validity, the births (fertility) and deaths (mortality, or survivorship) components of the cohort component model tend to be relatively uncontroversial. These components are relatively stable, at least over relatively short horizons, and the underlying drivers can be often captured with simple time series models in a way that, in our experience, end-users readily accept. Net migration, on the other hand, can be incredibly volatile, while being a quantitatively important component of national and - even more so – regional population change. In many cases it is difficult to project net migration accurately, and in a way that achieves high face validity.

At the national level, net migration consists of offsetting age-sex-specific immigration and emigration flows. A common approach adopted by many (multi)national statistical agencies, including the United Nations (2019), and taught in popular textbooks such as Rowland (2003), is to project a single value of overall net migration in each year for each projection 'scenario'. The projected path of net migration may be obtained by means of a time series model or simply set at plausible values. Disaggregation by age and sex can be based on an age-sex profile of historical net migration calculated by residual methods (with net migration equal to observed population change minus births plus deaths), with the age-sex profile being kept constant in projection years. This net migration approach remains popular despite of the well accepted critique that there is no such thing as a 'net migrant' (Rogers 1990).

An advance on the net migration approach is to obtain historical time-series of immigration and emigration, as advocated by Wilson and Rees (2021), which can then be projected separately using time series models and subsequently disaggregated by age and sex. A challenge to face validity of this approach may arise when the time series models do not include underlying drivers of international migration, in particular economic drivers and immigration policy changes.

Projections of migration at the subnational level add an additional layer of complexity. Overall (internal plus international) net migration approaches are relatively common at this level. A recent New Zealand example is Statistics New Zealand (2021). Projections with bi-directional migration flows (i.e. inflows and outflows) at the subnational level are rarer. However, in-flows and out-flows at the subnational level are comprised of a mixture of immigration, emigration, in-flows from other subnational areas within the country, and out-flows to other subnational areas within the country. To date, there have been few subnational population projection models developed that separately model all four types of migration (but see Poot et al. (2016) for a New Zealand approach). That creates a particular challenge for the face validity of subnational population projections. End-users may expect a model to consider differences in employment opportunities, incomes, house prices, amenity values, and access to services, as drivers of subnational migration flows. However, in conventional models based on net migration totals, as well as in models of directional migration flows, data on these drivers may not be known and are certainly hard to forecast. That can leave end-users disillusioned with what they perceive as unrealistic models of population change.

With respect to projecting directional internal migration flows, a relatively simple model – the gravity model – is readily available that has a long and distinguished history of use, as well as a remarkably good fit to observed migration flows (Poot et al. 2016). The gravity model of migration dates back to the early insights of Ravenstein (1885; 1889) and was further developed by Zipf (1946). The gravity model posits that the flow from region i to region j is positively associated with the 'economic size' of regions i and j (that can be proxied by their respective population sizes), and negatively associated with the distance between them. When relevant data are available the model can be augmented by the inclusion of a range of other variables that represent push and pull factors associated with migration (Lee 1966). However, this model can be improved upon by incorporating the interdependencies between the cells in the gross migration matrix, as formally introduced by Alonso (1978). Doubly-constrained gravity models such as introduced by Wilson (1971) are a special case.

Despite the attractiveness of the gravity model for modelling migration, including its empirical fit and ease of interpretation, incorporation of gravity into a subnational population projections model is rare. Cohen et al. (2008) and Kim and Cohen (2010) advocated the use of gravity models in the projection of international migration flows, but left as an open research question how well these models would work for internal migration projection. More recently, Husby and de Jong (2020) reported on a gravity model of internal migration flows that is incorporated into a municipality-level population projections model for the Netherlands. Their modelling framework was top-down, in the sense that the projections were constrained to sum to a national total, and hence they applied an 'origin-constrained' gravity model. They fit their model using geographically-weighted Poisson regression with fixed effects, and demonstrate good out-of-sample migration forecasts with their model. The internal migration flows are then included in their municipality-level population projections. A recent New Zealand application of the Alonso general modal of movements can be found in Cameron and Poot (2024).

In this paper, we report on a bottom-up subnational population projections model for 66 New Zealand territorial authorities (TAs). The model extends the conventional cohort component approach by projecting internal migration flows by means of the Alonso generalised gravity model. Regional immigration and emigration are also influenced by population size of a TA. To our knowledge, our model with gravity-based projections of migration flows is one of the first of its kind. We test the forecasting ability of the model by comparing out-of-sample projections of our model with those prepared by Statistics New Zealand (SNZ). SNZ projects the level of net migration by means of a single number for each TA. In terms of the standard criteria of mean absolute percentage error, our model performs better than SNZ's net migration-based one.

Additionally we demonstrate the efficacy of our approach by comparing our 2018-2048 projections with the official Statistics New Zealand (SNZ) 2018-base subnational population projections (SNZ 2021). We consider the impact of the dramatic reduction in international migration due to COVID-19 in our projections and the possible longer term effects of the 'reset' of immigration policies (NZPC 2021a). These modelling choices lead to differences between the two approaches in average population growth. However, we find that our projections are similar to the official projections in terms of the ranking of TAs by population growth rate.

We conclude that our methodology has several benefits over conventional approaches. The first is that it opens up for policy makers the 'black box' of net migration by explicitly considering international and internal migration at the regional level. This is essential in countries such as New Zealand due to the volatility of international migration and the large contribution of migration to regional population change. Secondly, we project inward and outward migration separately, which is important because their age-sex profiles are often very different. Thirdly, our methodology performs well in terms of out-of-sample forecasting of population change. Finally, the effects of major shocks and policy changes – such as those due to COVID-19 – are easier to integrate in our methodology than in the conventional net migration approaches. Overall, we argue that the gravity-model-based approach has greater face validity than conventional approaches.