UNIVERSITY OF WAIKATO

Hamilton
New Zealand

Using Land-Use Modelling to Statistically Downscale

Population Projections to Small Areas

Michael P. Cameron and William Cochrane

Department of Economics

Working Paper in Economics 15/12

November 2015

Corresponding Author

Michael Cameron
Economics Department
and NIDEA
University of Waikato
Private Bag 3105
Hamilton, 3240
NEW ZEALAND

Tel: +64 (07) 858 5082
Email: mcam@waikato.ac.nz

William Cochrane
Faculty of Arts & Social Sciences
and NIDEA
University of Waikato
Private Bag 3105
Hamilton, 3240
NEW ZEALAND

Email: billc@waikato.ac.nz

Abstract

Local government planners, property developers, large businesses and other stakeholders typically require good quality projections of the spatial distribution of the future population at the small-area level. Many approaches are available to project future populations, but all suffer from limitations due to their strict underlying assumptions or limited availability of data. In this paper we apply a novel approach to small-area population projection that combines cohort-component projections at the district level with grid-based land use projections at a fine (four-hectare) geographical scale. In our approach, residential population is directly estimated in the land use model, while a separate statistical model is used to link non-residential population to non-residential land use (by type). The model can then be used to project future small-area populations using projections of future land use from the land use model. We compare four data and model specifications for the statistical modelling, using either absolute land use area or principal components as explanatory variables, and using either OLS or Spatial Durbin model specifications. All four model combinations perform reasonably well for the Waikato Region of New Zealand, with good in-sample (2006) and out-of-sample (2013) properties. However, a naïve model based on constant shares of growth outperforms all four of our models in terms of forecast accuracy and bias. Notwithstanding the underperformance relative to a naïve model, our results suggest that land use modelling may still be useful, because the model is understandable by local authority planners and elected officials, and generates greater stakeholder 'buy-in' than black-box or naïve approaches.

Keywords

population projections; small-area projections; forecasting; land use

JEL Codes

C53, J11, Q56, R23, R52

Acknowledgements

1. Introduction

Local government planners, property developers, real estate agents, large businesses and other stakeholders need good projections or forecasts of the future spatial distribution of population for planning purposes (Foss 2002 and Isserman 1984). Indeed, Myers (2001, p.384) notes that 'planning analysts regard population statistics as integral to virtually all aspects of planning'. This includes planning for future land developments, schools, hospitals, child care centres, care services for elderly people, traffic flows, electorate boundaries, and so on. In the case of local government, this need for good data is often reinforced by legislation that increasingly calls for fiscal sustainability. For instance, the recently amended Local Government Act in New Zealand requires territorial and local authorities to engage in asset management planning with a fifty-year time horizon. In order to develop detailed asset management plans, local governments therefore need a good understanding of future population growth, not only in total but for particular localities within their districts or regions.

The risks to planners of planning on the basis of an inaccurate forecast of the future population distribution can be large. An overestimate of future population growth for a particular area will induce over-investment in infrastructure (such as roads, water and other utilities) with resultant costs on the local authority. These costs may be able to be passed onto private sector developers but, if not, they will be borne by local ratepayers. On the other hand, an underestimate of future population growth will lead to infrastructure being insufficient to meet the needs of the population, with costs (in terms of congestion or shortages of services) borne by the local populations. While accuracy is only one of several criteria on which small-area population forecasts might be judged (Tayman and Swanson 1996 and Tayman 2011), from the perspective of local authority planners the need for accurate and timely population forecasts is clear.

A range of government and non-government organisations (typically academics or consultants) produce population projections. Demographers and population modellers view population *projections* as distinct from population *forecasts*. A projection is a deterministically-driven path of future population growth, based on known assumptions about population change that are selected by the modeller. Different sets of assumptions will lead to different population projections. In contrast, a forecast is a 'most-likely' projection of the future, often based on the most likely set of assumptions for the model. Following Smith (1987) and Wilson and Rowe (2011), in this paper we note that the data are projections, but we evaluate them as if they were forecasts.

The projection assumptions that will lead to the most accurate forecast of future population are unknown (and some might argue, unknowable). This means that all forecasts of future population will be subject to error. The magnitude of forecast error tends to be substantially larger for areas with small populations, in comparison with more populous areas

(Cameron and Poot 2011, Smith 1987, Smith and Shahidullah 1995), which poses a particular problem for local authority planners interested in the distribution of population over small areas. Moreover, the methods for projecting population at small-area level are under-developed relative to projection methods for larger areas, although research in this area has increased recently (Chi 2009, Chi, Zhou and Voss 2011). In particular, it is unclear whether modellers should adopt a top-down approach to population projections (projecting large areas such as the national population first, followed by sequentially smaller component areas) or a bottom-up approach (projecting small areas first; then deriving projections for larger areas by summing the projections of their component small areas).

In this paper, we adopt a novel approach that uses a top-down cohort component population projection to define a district-level population, and then allocates the population spatially to small areas using statistical downscaling based on a model of land use. We compare four different model specifications based on land use, and then compare with two naïve projections. The remainder of the paper proceeds as follows. First, we describe the range of approaches that have been applied to the projection of population at small areas, and outline the strengths and weaknesses of each approach. Second, we describe the study area and the projection model we employ, as well as the method for evaluating the in-sample and out-of-sample forecasting performance of the model. Fourth, we present the results of the evaluation. Finally, we discuss the implications of our results for the use of similar models in local authority planning.

2. Small-Area Population Projections

There is no universally accepted definition of a 'small area' in terms of population projections (Smith and Morrison 2005). Wilson (2015a, p.335) notes that small areas tend to be 'spatial units at the detailed end of any geographical classification for which limited data exist'. This suggests that small areas would include sub-county-level areas such as census tracts in the U.S., or sub-district-level areas in other countries.

A variety of methods have been applied by demographers and population modellers to develop small-area population projections to satisfy planning needs. These methods can be generally categorised into four types:

(1) Naïve models, for example, extrapolation, or growth share models;
(2) The 'traditional' demographic cohort component model;
(3) Statistical methods, using data such as building consents, and
(4) Urban growth modelling approaches.

While there may be some overlap between these methods, and it is possible (and sometimes desirable) to combine approaches (indeed that is what we do in this paper), the following paragraphs discuss the relative strengths and weaknesses for local authority planning of each of these methods separately.

Naïve models are essentially simple extrapolations of past populations. The simplest of these models involve an assumption of no population change (i.e. constant population), linear population growth based on past growth trends, and exponential growth (i.e. constant population growth rates) based on past growth trends. Slightly more sophisticated are growth share models, which are top-down models where the population is initially projected at a higher geographical level, and then that projected growth is shared between the different small areas. For example, the total population for a state or county may be projected first, and then the growth share model used to allocate the state- or county-level population to the census tract level.

Wilson (2015a) tests a wide array of these naïve models for small areas in Australia, England/Wales, and New Zealand. He finds that, in terms of individual models, a constant share of population (CSP) model works best (smallest error) for England/Wales, and a variation on a constant share of growth (CSG+) model works best for Australia and New Zealand. Following White (1954), the CSG+ model assumes a constant share of growth for those areas that experienced positive population growth in the base period (that is, just before the start of the projection), and no growth for areas that declined in population in the base period.

Naïve models offer the advantage of simplicity of calculation, and have limited data requirements. The simplest (no growth, linear or exponential growth) models only require time series data on past population for each small area – even two data points of observed population will be enough to project the population. Growth share models require more detailed data and a separate model of population change operating at a higher geographical level. However, the necessary projections at the state, district, or national level are often publicly available from national statistical agencies.

Naïve models also tend to perform reasonably well, in terms of forecast accuracy, when compared with more sophisticated models (van der Gaag, van Wissen, Stillwell, and Kupiszewski 2003). Wilson (2015a) tested the forecast performance of simple models as well as averaged models (comprising the average of one or more of the simple models), and composite models (comprising different models for different types of areas). He found that the simple models performed nearly as well as the more sophisticated averaged or composite models, and the best simple models outperformed many of the possible permutations of averaged or composite models. Even within the class of naïve models, more sophisticated models are not necessarily always better. For instance, Rayer (2008) and Rayer and Smith (2010) have found linear extrapolation to be more accurate than growth share models for sub-county areas in Florida.

However, naïve models have practical limitations. First and foremost, they lack a strong theoretical basis. The purely mechanistic and deterministic application of past time trends into the future may be appropriate for small areas that have stable and predictable

growth paths, but most small areas are subject to unexpected changes in population. By ignoring the demographic or urban/land use drivers of population change (that the more sophisticated models outlined below feature), these models will fail to adequately account for changes in these drivers. On a related note, because local contextual factors are not incorporated into the model, naïve models of population change are difficult to justify to planners or, importantly, to elected officials. This is because there are no mechanisms for local policy to affect the future population distribution. This can lead to a lack of 'buy-in' from important end-users of the projections.

The traditional workhorse of demographic projections is the cohort component model (CCM), the modern application of which dates back to Whelpton (1928, 1936). In the CCM the population is projected by first projecting the three components of population change: (1) births, typically projected by means of age-specific fertility rates applied to women of childbearing ages; (2) deaths, typically projected by means of age-sex-specific mortality (or its complement, survivorship) rates applied to the population of each age and sex; and (3) migration, which may be projected in a number of different ways (van der Gaag et al 2003; Cameron and Poot 2011).

The advantage of the CCM is that all of the demographic drivers of population change are explicitly included in the model, because the CCM formula is an identity. Depending on the process used for modelling each component, other known drivers of population change can be explicitly included through their influences on fertility, mortality, and/or migration. Because the drivers are included, the CCM is typically intuitively understandable for local authority planners and elected officials. The key challenge for modellers is to demonstrate to these stakeholders how local policy influences are incorporated into the assumptions for future fertility, mortality, or especially migration. A second advantage of the CCM is that it generates the age and sex distributions of the population automatically, rather than simply projections of total population. Where this additional detail is required the CCM may be the best approach (Wilson and Rowe 2011).

However, at small-area levels, the CCM faces significant challenges. First, the data necessary for deriving assumptions about future fertility, mortality and migration may not be available at small geographical scales (Smith and Shahidullah 1995, Wilson 2015b). Where these data are available, data quality or precision may be so low that it may be difficult to derive robust age-specific rates for each component (fertility, mortality, inward and outward migration) at local levels (Wilson 2015b, Tayman, Schafer and Carter 1998). For instance, age-sex-specific data are typically required to estimate fertility and mortality rates. At the small-area level, the counts of births and deaths that occur each year (particularly when you consider age-specific counts) are small and highly variable, such that the estimation of fertility and mortality rates becomes extremely challenging.

Secondly, despite the promise that incorporating drivers of each component (fertility, mortality, migration) holds, most CCM models fail to adequately take account of a myriad of socio-economic, infrastructural, physical land use and other contextual factors that exert substantial influence over the spatial allocation of population and households at smaller geographical levels. Typically, these factors are excluded due to data unavailability and the inability to reliably forecast them. Contextual factors matter much more at the local level, such as the availability of suitable land, services and amenities, and the plans of public and private land developers (Murdock, Hamm, Voss, Fannin and Pecotte 1991). In particular, land use and availability constraints, planning constraints, and the availability of infrastructure are all variables that local authority planners would expect to impact on the future population distribution at the small-area level. Drivers of migration such as employment rates and expected incomes are known to be important determinants of migration flows. Unsurprisingly, the omission of these important drivers from CCM models means that there is only a limited role in the model for local policy, and reduces buy-in from local authority planners and elected officials.

Of course, the quantitative test of the CCM is whether it outperforms other models in evaluations of small-area projections. Unfortunately, past studies have shown that CCMs do not outperform simpler methods in projecting small area populations (Smith 1997, Smith and Tayman 2003 and Rayer 2008).

Statistical models offer one way to include the important contextual variables that are absent from the naïve models, and often missing from CCM models as well. For instance, regression models have become increasingly common in small-area projections (Alho and Spencer 2005), including more recently spatial regression models (Chi and Voss 2011, Chi, Zhou and Voss 2011). Spatial regression models may be preferred over aspatial models because the effects of the characteristics and contexts of neighbouring areas may also be important drivers of population change in each small area (Chi, Zhou and Voss 2011), and traditional regression models are unable to account for all of the spatial interactions (Anselin 1988; Harding, Vidyattama and Tanton 2011, Lesage and Pace 2009). Chi, Zhou and Voss (2011), for example, used spatial lag models to derive population projections for census tracts in Milwaukee, Wisconsin.

The advantage of statistical models is that the relationship between the drivers of population change and the resulting population projections is made explicit and transparent to local authority planners and other end-users of the projections. However, statistical methods also have limitations. Most important among these is that, because the statistical methods are typically purely data-driven and not based on an explicit demographic theoretical model, statistical models of small-area population often ignore the demographic drivers which are known to be causal in population change, that is, fertility, mortality, and migration. This limitation makes them open to criticism, from demographic modellers in particular. Secondly, as with the limitation for including more detailed population drivers within CCM models,

data availability may be a serious issue, and all data that is used within the statistical model must also be projected. For instance, if the statistical model uses building consent data, then projections of future building consents would be required in order to project the population (although in some models these projections might be obtained simultaneously).

Statistical models also suffer from a range of well-recognised issues. These include temporal instability of coefficients and over-fitting. Statistical models assume that coefficients are stable over time, but that need not be the case (Tayman and Schafer 1985). In fact, it is likely that coefficients change systematically over time as the size and spatial distribution of the population, and the characteristics of the population and the local environment, change. Overfitting occurs when the model is overly complex, containing too many parameters relative to the number of observations the model is based on. The consequence of overfitting is that random fluctuations in the data can be exaggerated, leading to inaccurate forecast performance.

Finally, like CCM models, statistical models have not been demonstrated to outperform even simple models of small-area populations in terms of forecast accuracy, even when the statistical model includes spatial interactions. For instance, Chi, Zhou and Voss (2011) find that their spatial lag model for Milwaukee does not unambiguously outperform projections derived from simple extrapolation methods.

The final category of small-area projection models is models based on urban growth modelling approaches, including:

(1) Cellular Automata (CA) modelling;
(2) Artificial neural networks;
(3) Fractal modelling;
(4) Agent-based modelling, and
(5) Decision-trees modelling.

CA modelling involves separating each area into a grid of cells, each of which has a number of characteristics (which may include population size). In each time step of the model, each cell may change its characteristics in response to shifts in the characteristics of neighbouring cells and changes in the nature of the system as a whole (see also the description of the land use model in the following section). For the sake of brevity, we do not explain in detail each of the remaining urban growth modelling methods, which are described and reviewed in Triantakonstantis and Mountrakis (2012). The advantages of urban growth modelling approaches include a much stronger theoretical base than statistical modelling, and that these models are able to more explicitly account for the local socio-economic conditions and physical and planning constraints at the small-area level.

However, the limitations of urban growth modelling approaches are similar to those for statistical models. The data requirements for urban growth models are very high, and a lot

of data may not be readily available at the small-area level. We note that this is particularly an issue for CA models, where grid cells are likely to be even smaller both geographically and in terms of population size than the small areas used in statistical models. Urban growth models also often do not specifically model the underlying demographic processes that drive population change, which is a significant limitation in the eyes of demographic modellers. However, this omission need not always be the case, and is certainly not true when the urban growth model is coupled with a demographic model, such as in our approach.

An alternative to applying one of the four approaches above is to combine two or more approaches in order to leverage their particular strengths, and attempt to address their limitations. One increasingly common combined approach involves using demographic projections such as CCM models to derive estimates of the future population at a relatively broad geographical scale, then using one of the other approaches to systematically downscale or apportion the population to the small-area level. Combining two approaches can take account of both the underlying demographic processes that drive population change, and the local-level conditions that primarily determine the spatial allocation of households and people (Wilson 2015b). Moreover, by combining two methods the demographic model is not overextended to a point where the data necessary to derive population projection assumptions (fertility, mortality, and migration) are not readily available.

Finally, the combined approach avoids one problem that is inherent in the solely top-down modelling approach. When the populations of larger areas are projected first, before the constituent small areas are projected, this leads to the tendency to assume that all constituent areas have similar growth trajectories to the larger geographic area (Tayman and Swanson 1996). Because the combined approach involves both top-down and bottom-up processes that are dependent but involve different methods, this tendency can be reduced.

In the combined approach the method of allocating population between small areas becomes the most important determinant of forecast accuracy at the small-area level. One method for achieving this allocation is to apportion population on the basis of constrained future housing availability (the 'housing unit' method). In this method, the future pattern of land use is first projected (whether by use of a formal modelling approach or simply by consulting future zoning plans), then the number of housing units that can be accommodated based on the pattern of land use is calculated. Finally, assumptions about the proportion of housing units that are occupied and the number of inhabitants per occupied housing unit are used to derive the population of each small area (Foss 2002). For instance, Roskruge, Cameron, and Cochrane, (2011) used future zoning changes, combined with assumptions about rates of infill housing, to project the theoretical maximum number of housing units in each CAU and to allocate population within the Waikato District of New Zealand. Overall, it appeared that the housing unit method worked well for urban areas that have a reasonably predictable spatial pattern of growth, but less well for slowly growing or declining areas.

Land use based models have been used to downscale or allocate population to small areas for at least the last two decades. Tayman (1996) reports results of a forecast based on a spatial interaction land use model for San Diego County. The land use model uses place-of-work employment to allocate population, such that the population tends to locate closer to their place-of-work, while constraining population based on each zone's capacity to accommodate additional residential development. Tayman and Swanson (1996) used similar models for San Diego and Dallas-Fort Worth. In a slightly more involved approach, Bell, Dean and Blake (2000) first used a land use model to disaggregate a CCM-based population projection for South Australia to the metropolitan district level for Adelaide, before employing a second model (called PUP) to allocate the district-level population to very small areas (less than a quarter of a square kilometre in size).

A final combined approach is to use spatial microsimulation to allocate population (Harding, Vidyattama and Tanton 2011). Spatial microsimulation involves using a model based on a spatially-explicit sample of the population, and transition parameters for demographic processes that allow future small-area populations to be derived. However, the application of spatial microsimulation to small-area population projections is still in its infancy, and this approach has yet to be rigorously evaluated in terms of forecast accuracy and bias.

3. Data and Methods

Data

Our particular application of small-area population projections is for the Waikato Region, in the North Island of New Zealand, using boundaries at the time of the 2013 Census of Population and Dwellings. The Waikato Region had a 2013 total population of approximately 425,000 (about 10 percent of the total New Zealand population). It has a central main city (Hamilton City) with a 2013 population of approximately 150,000, two districts that are peri-urban (Waikato District and Waipa District), and a number of other Territorial Authority (TA) areas[1] in whole or in part (refer to Table 1).[2]

As noted in the previous section, there is no commonly-applied definition of a small area for projections purposes. In this paper, we develop projections at the Area Unit (AU) level. Area Units are the next smaller geographical area below TAs in the geographical hierarchy used by Statistics New Zealand. They serve no particular administrative purpose – however, each AU is a distinct geographical entity, and in urban areas they generally coincide with suburbs and have a population of 3,000-5,000. As shown in Table 1, in the Waikato

---

[1] Territorial Authorities are the second tier of local government administration in New Zealand.

[2] The region is not a simple aggregation of the TAs because the region is largely based on a water catchment area, whereas the TA boundaries reflect administrative divisions that are historical and somewhat arbitrary.

Region there are 197 non-marine non-island AUs, with a mean population size in 2013 of 2,124 (median 1,840), and a range from a minimum of zero to a maximum of 7,750.

Table 1: Territorial Authority Populations for the Waikato Region, 2013

Territorial Authority	Population	Count of Area Units (AUs)	Mean AU Population	Median AU Population	Minimum AU Population	Maximum AU Population
Thames-Coromandel District	27,040	10	2704	2845	730	4490
Hauraki District	18,740	8	2343	1945	500	4790
Waikato District	64,890	31	2093	1860	0	5550
Matamata-Piako District	32,200	13	2477	2510	300	4520
Hamilton City	150,250	46	3266	3305	160	7750
Waipa District	46,380	29	1599	1300	200	3770
Otorohanga District	9,330	5	1866	1750	350	4180
South Waikato District	22,530	16	1408	1060	160	3690
Waitomo District (part)	9,330	7	1333	1000	210	4670
Taupo District (part)	34,120	28	1219	625	10	4410
Rotorua District (part)	3,640	4	910	870	160	1740
Waikato Region (Total)	418,450	197	2124	1840	0	7750

Statistical Downscaling Method

In this paper, we use a combined approach similar to those discussed in Section 1, which involves using statistical downscaling combined with projections of future land use to allocate projected TA-level populations to each AU. The three-step approach we adopted is

similar to that employed by Tayman and Swanson (1996) and Tayman, Schafer and Carter (1998), but uses a combined statistical and urban growth modelling approach to allocate population to small areas.

First, the population was projected at the TA level for the region (including for each part-TA) by the National Institute of Demographic and Economic Analysis, using a cohort component model (Cameron and Cochrane 2014). The 'Baseline Medium' TA-level projected populations were used as an input in the following stages, including a backcast projection from 2013 (the base year of the TA-level projections) to 2006.

Second, land use was projected using the Waikato Integrated Scenario Explorer (WISE) model. The WISE model is a systems-based integrated model that incorporates economic, demographic, and environmental components across the entire Waikato Region (Rutledge *et al.* 2008, 2010). The WISE model begins with a base land use map in 2006, incorporating 24 different land uses, of which there are three residential land use classes (medium-high density, low density, and lifestyle blocks) (Rutledge *et al.* 2010). At each annual time step, the economic and demographic sub-models generate demands for economic and residential land use, which are inputs into a dynamic, spatially explicit land use change model (Huser *et al.* 2009, van Delden, Gutierrez, van Vliet and Hurkens 2008). The demographic inputs into the WISE model are the TA-level population projections for the Waikato region developed in the first step.

The land use change model is a CA model specified at the level of four-hectare grid cells (200m x 200m). The CA model apportions land to different uses at each annual time step based on a combination of four factors:

(1) Zoning (which constrains the land uses that are available in each area);
(2) Suitability (the biophysical suitability of land for different uses);
(3) Accessibility (assesses the attractiveness of a location for different land uses based on proximity to desirable or undesirable features) and
(4) Local influence (assesses the attractiveness of a location for a land use based on the composition of land use in the surrounding neighbourhood).

The CA land use model attempts to meet the external demands for land (from the economic and demographic models) by assigning cells with the highest transition potentials (determined by their zoning, suitability, accessibility and local influence) to new land uses. Transitions are made at each annual time step.

The demand for residential land of each type is determined by first assigning a given proportion of population in each territorial authority to each residential land use type, and the residual proportion is spread across all non-residential land uses. The proportions are generally stable but vary over time for some TAs. Next, the number of residential land use cells of each type required is determined by combining the population in each residential land

use calculated in the first step with population density values for each residential land use type. These population densities also vary over time, between pre-determined maximum and minimum values. The area of each land use type (in hectares) and the residential population densities (by residential land use type) were exported from the WISE model for 2006 and 2013 for use in the next step.

In the third step, land use was used to statistically downscale the TA-level population projections to the AU level. This was achieved in two stages, projecting: (1) the population located in residential land uses and (2) the population located in non-residential land uses. In the first stage, the number of hectares of each residential land use type in each AU and the residential population densities (both from the WISE model) were used to calculate the residential population of each AU (that is, the population located in residential land uses) for each year (2006 and 2013). The difference between the sum of the residential populations across all AUs in each TA and the overall projected TA-level population provides an estimate of the total non-residential population in that TA (that is, the population located in non-residential land uses).

To estimate the non-residential population in each AU, linear regression models were used, with the 2006 TA-level non-residential population as the dependent variable, and the 2006 baseline non-residential land use (by type) as explanatory variables. That is, we estimated a regression model of the general form:

$$NRP_i = \alpha + N_{ki}\beta + \varepsilon_i \tag{1}$$

where NRP_i is the non-residential population of area unit i, N_{ki} is a vector of land uses k in AU i, and ε_i is an idiosyncratic error term. We test and report four alternative model and data specifications for this model: (I) standard ordinary least squares (OLS) regression, based on absolute land use; (II) standard OLS, based on principal components of land use; (III) a spatial Durbin model, based on absolute land use; and (IV) a spatial Durbin model, based on principal components of land use.

The rationales for testing these different specifications are as follows. Absolute land use (in hectares) is the most basic land use variable, and we tested this data specification because the coefficients for each land use type would be easy for planners to interpret (as the number of people per hectare). In contrast, principal components takes the land use dataset and converts it into a set of linearly uncorrelated components (Joliffe 2010). This avoids any problems of multicollinearity between the land use variables. We argue that the principal component specifications also allow the model to account for different types of internal land use structures that may be reflected in different land-use-specific population densities at the AU level.

Spatial Durbin models account for neighbourhood effects (that is, where the non-residential population in AU i is affected by the size of the non-residential population in

surrounding AUs) and for lag effects (where the non-residential population in AU i is affected not only by the amount of each land use type in that area unit, but also the amount of each land use type in surrounding AUs) (Anselin 1988, Lesage and Pace 2009). Spatial models for small-area population projections have been tried in the past, and found to underperform relative to simpler methodologies (Chi, Zhou and Voss 2011; Chi and Voss 2011). However, in our model only the non-residential population is projected using the spatial regression model, with the residential population having been determined previously. Thus, we might expect an improvement in modelling performance.

Eleven land uses were initially excluded from the models (bare surfaces; indigenous vegetation; other exotic vegetation; wetlands; fresh water; marine; aquaculture; utilities; mines and quarries; urban parks; and airports), because they were unlikely to contain much population. The three residential land uses were also excluded from the models, as the population in those land uses was already accounted for. That leaves ten land use variables in the model. Separate regression models were fitted for Waikato District, Hamilton City, and Waipa District, with a fourth regression model fitted for the remaining TAs (due to small individual sample sizes). The fourth model initially included TA-level fixed effects to account for unobserved differences in population density profile between each TA. Each regression model was reduced to a final preferred model by removing the least significant variable in a backward stepwise fashion until the root mean squared error (RMSE) was minimised. The resulting regression models are a reasonably good fit for the data, with adjusted coefficients of determination (R^2) between 0.17 and 0.80.[3]

The regression model coefficient estimates were then used, along with projected land use from the WISE model for 2013, to provide a projection of the non-residential population of each AU in 2013. When added to the residential population from the first stage of step 3, the sum provides an un-scaled population projection for each AU. However, two issues arose with these un-scaled projections: (1) the projections demonstrated significant discontinuity with the known population trend between 2001 and 2006 for a number of CAUs; and (2) a number of CAUs were projected to quickly fall to zero (or negative) population. To reduce the impact of the discontinuities, the in-sample residual was calculated for each CAU in 2006 (being the difference between the actual 2006 population and the estimated 2006 population). This in-sample residual for each AU was added to the projected AU population. This reflects the fact that the residuals in the population projection model are likely to be correlated over time.

To reduce the impact of projected de-population of (particularly rural) AUs, each un-scaled AU population projection was constrained so that population would not fall by more than 25 percent over a ten-year period. This maximum constraint is similar to the maximum long-run population decline observed in any AU over the period 1996-2006. Moreover, this

---

[3] We omit detailed results for the sake of brevity, but these are available on request from the authors.

adjustment is justifiable because the spatial distribution of population is subject to a substantial degree of inertia – once houses have been constructed in a given location, some population is likely to remain in that location for a long time. That is, population decline at small spatial scales is a relatively slow process, unlike that projected in the initial unconstrained models.

Finally, the combined population of all AUs in each TA was constrained to be consistent with the projected population of the TA from the cohort component model. Discrepancies between the AU-based population total and the TA-level projection were eliminated by applying a common scaling factor to the AU populations for each TA, calculated as the ratio of the projected TA-level population to the sum of the unconstrained AU populations.

Evaluation Method

We evaluate the performance of our approach in two ways. First, we evaluate the in-sample performance of the model. Specifically, we compare the four alternative regression models for projecting non-residential population (and the resulting estimates of the AU-level populations) in 2006, by comparing the estimated with actual populations. Second, we evaluate the out-of-sample forecast accuracy by doing a post-hoc comparison of the small area forecasts with data from the 2013 estimated usually resident populations (based on the 2013 Census). We compare the forecast accuracy of our four models with that of two naïve models: (1) a linear extrapolation, that takes the population change from 2001 to 2006, and extrapolates this to 2013; and (2) a CSG+ (modified constant share of growth) model, which assumes a constant share of TA growth for each AU that experienced positive population growth between 2001 and 2006, and no growth for areas that declined in population between 2001 and 2006 (White, 1954; Wilson, 2015a). Rather than using a population projection model in the CSG+ model, we use the change between actual 2006 and 2013 estimated usually resident populations. This provides an over-conservative estimate of the degree of error and bias in the CSG+ model.

We make use of multiple measures of forecast error and bias. Following Wilson (2015b), our primary measure of forecast accuracy is weighted mean absolute percentage error (WMAPE). This measure is a weighted mean of the absolute percentage errors, with the weights being the size of the actual populations in the year projected (Siegel 2002). WMAPE is preferable to other measures (such as Mean Absolute Percentage Error) when there is a wide range of population sizes. The AU populations in our study area range from zero to 7,750 in 2006, which makes WMAPE the most suitable measure.

We also report the median absolute percentage error (MedAPE), the median algebraic percentage error (MedALPE), and the root mean square error (RMSE). MedAPE and RMSE both measure forecast precision because the direction of the error does not affect these

measures, while MedALPE measures forecast bias. Although Mean Absolute Percentage Error (MAPE) and Mean Algebraic Percentage Error (MALPE) are the most commonly used measures of forecast accuracy and bias respectively (Tayman 1996), MedAPE and MedALPE are preferable over MAPE and MALPE. This is because using the median error reduces the impacts of extreme outliers (that is, unusually large, or small, errors) and the skewed nature of the distribution of error in small populations, on the overall measures of error and bias (Tayman and Swanson 1999). For instance, Tayman (1996) shows that MAPE tends to overstate the error, and that the degree of overstatement is largest for areas with the smallest population size. In contrast to these other measures, RMSE penalises the forecaster for forecasts that are further from the actual population (Stoto 1983), which may be helpful for risk averse planners adopting a minimax approach, that is, where forecasts will provide the planner with greater utility if the largest errors are minimised.

4. Results

Table 2 shows the results of the evaluation of the in-sample (2006) and out-of-sample (2013) performance of the four model and data specifications (I-IV), using all four error measures (WMAPE, MedAPE, MedALPE, and RMSE). We include two out-of-sample comparisons: (1) using the raw statistical model described in the previous section; and (2) using the statistical model, but carrying forward the 2006 in-sample residual and using it to modify the 2013 projection.

Table 2: In-Sample and Out-of-Sample Model Performance

Error Measure	Model I	Model II	Model III	Model IV
In-sample				
WMAPE (%)	19.0	19.0	16.2	16.3
MedAPE (%)	17.5	17.6	14.8	14.3
MedALPE (%)	-2.7	-2.9	-0.7	-0.1
RMSE (%)[*]	26.6	26.5	23.3	22.8
Out-of-sample				
WMAPE (%)	20.8	20.5	19.0	18.2
MedAPE (%)	19.3	19.1	16.2	16.5
MedALPE (%)	-0.6	-0.6	-1.9	2.1
RMSE (%)[*]	28.3	28.4	27.0	25.9
Modified Out-of-sample[**]				
WMAPE (%)	7.5	6.7	9.0	6.7
MedAPE (%)	5.7	4.8	6.7	4.8
MedALPE (%)	-0.8	0.0	-1.4	0.0
RMSE (%)[*]	14.3	14.0	15.8	14.0

[*] As a percentage of the mean AU population.

[**] Modified out-of-sample measures include a correction, whereby the in-sample residual is carried forward to form part of the forecast.

Overall, the models exhibit a moderate degree of accuracy, with in-sample errors of between 14.3 and 19.0 percent. There is an overall downward bias in the models, as the MedALPE values are consistently negative. Unlike Wilson and Rowe (2011) our measures of WMAPE are larger than MedAPE, which probably reflects that absolute errors are largest for the AUs with larger populations. In terms of in-sample performance, Models I and II perform similarly, but are clearly dominated by Models III and IV (the spatial regression models), which exhibit smaller degrees of both error and bias. Comparing Models III and IV, both are similar in terms of error, but Model IV exhibits a smaller degree of bias. The median extent of bias in Model IV is a 0.1 percent under-projection, compared with a 0.7 percent under-projection for Model III. The comparison between the four models is similar in the two out-of-sample comparisons.

Comparing the in-sample with the first out-of-sample results demonstrates that the land-use-based projection model performs nearly as well at seven years after baseline as it does in the baseline year. There is little degradation of performance over time for any of the models, with WMAPE increasing by between 1.5 and 2.8 percentage points between the in-sample and out-of-sample measures.

In the modified out-of-sample measures, Model III is clearly worse than other models, and Models II and IV perform best and nearly identically. There is no evidence of forecast bias in either of these models, and WMAPE is just 6.7 percent. Comparing the out-of-sample and modified out-of-sample results demonstrates the substantial performance improvement that is obtained by carrying forward the in-sample residual. Across all models this reduces the WMAPE by between one half and two thirds.

Table 3 shows the results of the out-of-sample comparison between our four models and the two naïve models (linear extrapolation, and CSG+). To ensure comparability between the naïve models and our models, we use our models to allocate the 2013 estimated usually resident population, rather than the 2006-base projected TA-level populations for 2013. Thus the error measures differ slightly from those reported in Table 2. As with the last comparison in Table 2, Models II and IV perform the best of our four models. They also perform better than the naïve linear extrapolation on three out of the four error measures (i.e. all except RMSE). However, the CSG+ model performs the best on all error measures, with a WMAPE of 5.6 percent, 1.1 percentage points better than Models II and IV.

Table 3: Comparative Out-of-Sample Model Performance

Error Measure	Model I	Model II	Model III	Model IV	Linear	CSG+
WMAPE (%)	7.3	6.7	8.7	6.7	7.6	5.6
MedAPE (%)	5.6	5.0	6.4	5.0	6.4	4.5
MedALPE (%)	-1.7	-1.7	-0.8	-1.7	-2.3	-0.3
RMSE (%)[*]	14.1	13.6	15.6	13.6	12.6	10.3

5. Discussion and Conclusions

In this paper we report the results of new land-use-based models for small area population projections, and compare those models' forecast accuracy with that of naïve projections based on linear extrapolation and a modified constant-shares-of-growth model. Our approach can readily be employed to projections in other areas, but necessarily requires a land use model. However, the land use model need not be as detailed as that employed here.

To date, few studies have compared the forecast accuracy of combined projection models with that of simpler models. Our preferred model (Model II) uses the WISE land use model to derive the residential population of each AU, and a spatial Durbin model using absolute non-residential land use (in hectares) as explanatory variables to derive the non-residential population of each AU. We prefer Model II over the similar Model IV because of the ease of interpretation of coefficients for end-users. This preferred model has an out-of-sample WMAPE of 6.7 percent over a seven-year projection horizon.

The error in our preferred model compares favourably with previous studies that use a variety of models (and measures of error). Wilson and Rowe (2011) found WMAPEs after five years varied from 6.0-7.3 percent for areas with 2,000-4,999 population, and 5.0-6.7 percent for populations of 5,000-14,999, for projections of the population of Queensland, Australia. These WMAPEs increased to 8.2-11.3 percent and 6.7-9.5 percent respectively for a ten-year projection horizon. Tayman, Schafer and Carter (1998) report that MAPE is a decreasing function of population size, based on San Diego data and a projection model that uses land use to allocate populations to small areas. They show MAPEs for a ten-year projection that range from 72 percent for populations of 500, to 39 percent for populations of 5,000, and to 10.5 percent for populations of 50,000. Using the same model, Tayman (1996) found a MAPE of 21 percent and MedAPE of 15 percent in ten-year projections of census tracts in San Diego County. Smith and Shahidullah (1995) calculate a MAPE of 19 percent for a ten-year projection of census tracts in Dallas-Fort Worth. Tayman and Swanson (1996), using a land-use-based model for census tracts in Detroit, Dallas-Fort Worth, and San Diego, found MAPEs of between 18.6 and 28.5 percent for a 10-year forecast horizon. For comparison, the unweighted out-of-sample MAPE for our preferred Model II is 10.9 percent for a seven-year projection, which is substantially lower than those reported in these previous studies. We further note that census tracts typically have larger population sizes than the area units we project, which further demonstrates the efficacy of our approach.

The source of the improved projections performance of our models, relative to previous projections models, is predominantly generated by the carrying-forward of in-sample residuals. Without carrying forward the residuals, the out-of-sample performance of our models looks much more similar to those of other models. This procedure makes sense for statistical and urban growth models (but not for extrapolation or CCM models). If other

studies using statistical models, such as Chi, Zhou and Voss (2011), carried forward residuals in their forecasts, then their model performance may look much better.

We also demonstrate that land-use-based population models that account for spatial interdependence (Models II and IV) outperform models that ignore these effects (Models I and III). The population density of a given small area reflects a complex interplay of the land use of that particular area, and the land uses of surrounding areas. For instance, urban land uses will have quite different characteristics and population densities than rural land uses, even within the same category of land use. Spatial Durbin models allow us to capture the spatial dependence, and small area population projection models should make more use of spatial models, as advocated by Chi and Voss (2011).

However, despite their good performance in comparison with past modelling efforts, our forecasts do not outperform the naïve CSG+ model (White 1954; Wilson 2015a). The inability of complex models to outperform simple models in projections of small area population is a general finding in the literature on small area population projections (van der Gaag et al 2003). However, we do note that the WISE land use model we employed in our analysis has recently undergone significant improvement, with input from a wide group of local authority planners. This improved model, which operates with a 2013-base land use map, demonstrates substantially better in-sample performance, with a nearly 30 percent reduction in WMAPE (based on our Model I) to 13.5 percent. Unfortunately, out-of-sample model testing based on this new land use model will not be possible until after the 2018 Census, but these initial results are extremely promising.

There are a number of limitations to the models presented here. First, the projections do not include an explicit measure of uncertainty. Instead, we simply forecast the AU populations as point estimates. However, the measures of forecast error that we calculated in our paper could be used to estimate uncertainty (Tayman 2011). It is worth noting that the degree of uncertainty present in population projections at smaller geographic levels is substantially larger for smaller populations (Cameron and Poot 2011), so understanding better the uncertainty in our estimates is clearly important. Secondly, because the forecasts are based on a statistical model at the small-area level, they potentially suffer from the same limitations as statistical models outlined in the introduction. However, because the statistical model is only used to project the non-residential population, rather than the whole population, these problems are somewhat mitigated. Thirdly, we evaluate our projections based on only a single period in time and a single region of New Zealand. It may be that demographic trends fit our model particularly well (or not so well) by chance alone. The model will be evaluated further in later periods, but should also be applied to other regions and contexts.

Fourthly, it is likely that small area population projections present a problem of endogeneity. If projections are used in planning decisions then they may become somewhat self-fulfilling prophecies. For instance, if population is projected to increase in a given AU,

then planners may create infrastructure that supports the expected additional population, leading to more development in that AU and consequently more population. However, if population had been projected to increase elsewhere instead, then infrastructure spending, development and population growth would be directed towards that other area instead. Thus, small area population projections should be used as one tool among many in the planning process.

Finally, despite the forecast accuracy of the models we presented in this paper being lower than naïve models, the land-use-based models do serve an important purpose. Too often, population projection models are seen by local authority planners and elected officials as 'black boxes' or academic curiosities that have little relevance to the real world. As Rainford and Masser (1987) note, bridging the gap between the technical aspects of forecasting and the needs of planners is both important and difficult. Achieving 'buy-in' from planners and elected officials is imperative in ensuring that population projections are understood and used effectively to achieve improved planning outcomes. One part of this is to ensure that planners can recognise that planning and policy have demonstrable effects on the projected populations at the small-area level. Our models have been very successful in this, and are being used extensively in long term planning processes at the local and regional level. Further enhancements to the model, including the improved land use modelling described above, will likely further increase the acceptance of planners for integrated modelling approaches.

References

Alho, J. M. and Spencer, B. D. (2005). *Statistical Demography and Forecasting*, Springer, New York.

Anselin, L. (1988). *Spatial econometrics: Methods and models,* Kluwer Academic Publishers, Dordrecht.

Bell, M., Dean, C., and Blake, M. (2000). Forecasting the pattern of urban growth with PUP: a web-based model interfaced with GIS and 3D animation, *Computers, Environment and Urban Systems,* 24, 559-581.

Cameron, M. P. and Cochrane, W. (2014). Population, Family and Household, and Labour Force Projections for the Waikato Region, 2013-2063, research report commissioned by the Waikato Regional Council Hamilton, New Zealand: University of Waikato

Cameron, M. P., and Poot, J. (2011). Lessons from stochastic small-area population projections: The case of Waikato subregions in New Zealand, *Journal of Population Research, 28*(2-3), 245-265.

Chi, G. (2009). Can knowledge improve population forecasts at subcounty level? *Demography, 46*(2), 405-427.

Chi, G. and Voss, P. R. (2011). Small-area population forecasting: Borrowing strength across space and time, *Population Space and Place, 17*(5), 505-520 .

Chi, G. Zhou, S. and Voss, P. R. (2011). Small-area population forecasting in an urban setting: A spatial regression approach, *Journal of Population Research, 28*, 185-201.

Foss, W. (2002). Small area population forecasting, *The Appraisal Journal, 70*(2), 163-172.

Harding, A., Vidyattama, Y., and Tanton, R. (2011). Demographic change and the needs-based planning of government services: Projecting small area populations using spatial microsimulation, *Journal of Population Research,* 28, 203-224.

Huser, B., Rutledge, D., van Delden, H., Wedderburn, M. E., Cameron, M., Elliot, S., Fenton, T., Hurkens, J., McBride, G., McDonald, G., O'Connor, M., Phyn, D., Poot, J., Price, R., Small, B., Tait, A., Vanhout, R., Woods, R. A. (2009). Development of an integrated spatial decision support system ISDSS for local government in New Zealand, Proceedings of the MODSIM09 International Congress on Modelling and Simulation pp.2370-2376, Lincoln University, Christchurch.

Isserman, A. M. (1984). Projection forecast and plan: On the future of population forecasting, *Journal of the American Planning Association,* 50, 208-221.

Joliffe, I. T. (2010). *Principal component analysis,* 2nd edition, Springer, New York.

LeSage, J. and Pace, R. K. (2009). Introduction to spatial econometrics, CRC Press, Boca Rato, FL.

Murdock, S. H., Hamm, R. R., Voss, P. R., Fannin, D. and Pecotte, B. (1991), Evaluating small-area population projections, *Journal of the American Planning Association,* 57 (4), 432-443.

Rainford, P. and Masser, I. (1987), Population forecasting and urban planning practice, *Environment and Planning A, 19*(11), 1463-1475.

Rayer, S. (2008). Population forecast errors: A primer for planners, *Journal of Planning Education and Research,* 27, 417-430.

Rayer, S., and Smith, S. K. (2010). Factors affecting the accuracu of subcounty population forecasts, *Journal of Planning Education and Research,* 30, 147-161.

Roskruge, M., Cameron, M. P. and Cochrane, W. (2011). Waikato District Sub-district Population Projections 2011-2031, research report commissioned by Waikato District Council Hamilton: National Institute for Demographic and Economic Analysis, University of Waikato.

Rutledge, D. T., Cameron, M., Elliott, S., Fenton, T., Huser, B., McBride, G., McDonald, G., O'Connor, M., Phyn, D., Poot, J., Price, R., Scrimgeour, F., Small, B., Tait,A., van Delden, H., Wedderburn, L., Woods, R. A. (2008). Choosing regional futures: challenges and choices in building integrated models to support long-term regional planning in New Zealand, *Regional Science Policy and Practice, 1*(1), 85-108.

Rutledge, D., Cameron, M., Elliott, S., Hurkens, J., MacDonald, G., McBride, G., Phyn, D., Poot, J., Price, R., Schmidt, J., van Delden, H., Tait, A. and Woods, R. (2010). WISE – Waikato Integrated Scenario Explorer, Technical Specifications Version 1.1 Research Report commissioned by Environment Waikato Hamilton: Landcare Research.

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

Smith, S. K. (1987). Tests of forecast accuracy and bias for county population projections, *Journal of the American Statistical Association,* 82, 991-1003.

Smith, S. (1997), Further thoughts on simplicity and complexity in population projection models, *Journal of the Forecasting,* 13, 557-565.

Smith, S. K. and Morrison, P. A. (2005). Small area and business demography, in *Handbook of Population* Eds D Poston, M Micklin New York: Springer. pp.761-785.

Smith, S. K. and Shahidullah, M. (1995). An evaluation of population projection errors for census tracts, *Journal of the American Statistical Association, 90*(429), 64-71.

Smith, S. K., and Tayman, J. (2003). An evaluation of population projections by age, *Demography, 40*(4), 741-757.

Stoto, M. (1983). The accuracy of population projections, *Journal of the American Statistical Association*, 78, 13–20.

Tayman, J. (1996). The accuracy of small-area population forecasts based on a spatial interaction land-use modeling system, *Journal of the American Planning Association, 62*(1), 85-98.

Tayman, J. (2011), Assessing uncertainty in small area forecasts: State of the practice and implementation strategy, *Population Research and Policy Review,* 30, 781-800.

Tayman, J. and Schafer, E. (1985). The impact of coefficient drift and measurement error on the accuracy of ratio-correlation population estimates. *Review of Regional Studies,* 15, 3-10.

Tayman, J. Schafer, E. and Carter, L. (1998). The role of population size in the determination and prediction of population forecast errors: An evaluation using confidence intervals for subcounty areas, *Population Research and Policy Review, 17*(1), 1-20.

Tayman, J., and Swanson, D. A. (1996). On the utility of population forecasts, *Demography, 33*(4), 523-528.

Tayman, J. and Swanson, D. A. (1999). On the validity of MAPE as a measure of population forecasting accuracy, *Population Research and Policy Review,* 18, 299-322.

Triantakonstantis, D., and Mountrakis, G. (2012). Urban growth prediction: A review of computational models and human perceptions, *Journal of Geographic Information System, 4*(6), 555-587.

van Delden, H., Gutiérrez, E., Van Vliet, J., and Hurkens, J. (2008). Xplorah a multi-scale integrated land use model, in proceedings of the iEMSs Fourth Biennial Meeting: Integrating Sciences and Information Technology for Environmental Assessment and Decision Making Eds M Sànchez-Marrè, J Béjar, J Comas, A Rizzoli, G Guariso International Environmental Modelling and Software Society.

Van der Gaag, N., van Wissen, L., Rees, P., Stillwell, J. and Kupiszewski, M. (2003). *Study of Past and Future Interregional Migration Trends and Patterns within European Union Countries: In Search for a Generally Applicable Explanatory Model,* Netherlands Interdisciplinary Demographic Institute, The Hague.

Whelpton, P. K. (1928). Population of the United States 1925 to 1975, *American Journal of Sociology, 34*(2), 253-270.

Whelpton ,P. K. (1936). An empirical method for calculating future populations, *Journal of the American Statistical Association,* 31, 457-473.

White, H. R. (1954). Empirical study of the accuracy of selected methods of projecting state populations, *Journal of the American Statistical Association,* 29, 480-498.

Wilson, T. (2015a). New evaluations of simple models for small area population forecasts, *Population Space and Place,* 21, 335-353.

Wilson, T. (2015b). Short-term forecast error of Australian local government area population projections, *Australasian Journal of Regional Studies, 21*(2), 253-275.

Wilson, T., and Rowe, F. (2011). The forecast accuracy of local government area population projections: A case study of Queesnland, *Australasian Journal of Regional Studies, 17*(2), 204-243.