Author/s:
Wilson, T;Grossman, I

Title:
Evaluating Alternative Implementations of the Hamilton-Perry Model for Small Area Population Forecasts: the Case of Australia

Date:
2022-01-05

Citation:
Wilson, T. & Grossman, I. (2022). Evaluating Alternative Implementations of the Hamilton-Perry Model for Small Area Population Forecasts: the Case of Australia. SPATIAL DEMOGRAPHY, 10 (1), pp.1-31. https://doi.org/10.1007/s40980-021-00103-9.

Persistent Link:
https://hdl.handle.net/11343/332348

Evaluating alternative implementations of the Hamilton-Perry model for small area population forecasts: the case of Australia

Tom Wilson*

Irina Grossman‡

* Corresponding author. Address: Melbourne School of Population and Global Health, The University of Melbourne, 207 Bouverie St, Melbourne, Vic 3010, Australia

Email: wilson.t1@unimelb.edu.au

ORCID: 0000-0001-8812-7556

‡ Melbourne School of Population and Global Health, The University of Melbourne

Email: irina.grossman@unimelb.edu.au

ORCID: 0000-0002-5761-6194

Acknowledgements

We thank the anonymous reviewers for their helpful comments on the initial version of this paper.

Abstract

Small area population forecasts are widely used across the public and private sectors, with many users requiring forecasts broken down by sex and age group. The preparation of small area age-sex population forecasts across a whole country or State with a multiregional cohort-component model is usually a time-consuming and expensive task. It involves the purchase of large datasets, considerable amounts of complex data preparation and assumption-setting, and substantial amounts of staff time. A quicker and lower-cost alternative is to use a reduced form cohort projection model, such as the Hamilton-Perry model. This paper presents an evaluation of various implementations of the Hamilton-Perry model, including an alternative version employing a combination of Cohort Change Ratios and Cohort Change Differences. It also evaluates the effects on forecast accuracy of smoothing the age profiles of Cohort Change Ratios and Differences, and constraining to independent population forecasts. Population 'forecasts' were created for small areas of Australia over the horizon 2006-16 and compared against population estimates. The most accurate implementation is found to be the Hamilton-Perry model using a combination of Cohort Change Ratios and Cohort Change Differences, smoothed age profiles, and with constraining to independent forecasts.

Key words

Small area; population forecasts; Hamilton-Perry model; Australia

1. Introduction

Small area population forecasts are regularly used across the public and private sectors to inform various planning and investment activities, such as childcare provision, school infrastructure development, market assessment, labour force needs, the adjustment of electoral districts, transport usage, and water and power demand. Many users require population forecasts by age group and sex. The preparation of small area population forecasts by age and sex across a whole State or country is usually a substantial undertaking. It can involve many thousands of areas, the purchase of a large amount of input data at considerable cost, and the dedication of a lot of time and effort preparing all the data inputs, programming a cohort-component model, setting forecast assumptions, and validating outputs. The project can take many person-months of time. Methods and tools which can reduce the time, cost, and complexity of this work are potentially very useful.

One approach to simplifying the task, which has proved popular over the last decade or so, is to employ the Hamilton-Perry projection model (e.g., Baker et al. 2017; Baker et al. 2020; Hauer 2019; Inoue 2017; Swanson et al. 2010; Tayman et al. 2021). This is a reduced version of the cohort-component model which retains the cohort aspect of that model, but omits the component part (Hamilton and Perry 1962). Instead of each cohort's population being subject to change from mortality and inward and outward migration as in the cohort-component model, cohort population change is projected directly. A cohort's population in five years' time (and five years older) is projected simply as the cohort's population at the present multiplied by a Cohort Change Ratio (CCR). The key advantage of the model is that age-sex population forecasts can be prepared relatively quickly and easily, and with little data. The minimum data requirements comprise just two sets of population estimates by sex and age group for each small area, one set for the jump-off year of the forecasts and the other for five (or in some cases ten) years earlier, which are used to calculate Cohort Change Ratios. This represents a huge saving in time and cost relative to the cohort-component model. However, the main trade-off is that there are no demographic components of change. Population is not projected as the result of individual fertility, mortality and migration processes; forecast assumptions cannot be formulated for each of these demographic processes; and there are no projected components of change in the outputs. Nonetheless, the Hamilton-Perry model permits the creation of age-sex population forecasts in situations where data limitations prevent the application of a regular cohort-component model.

An important question is then: how accurate are forecasts produced by the Hamilton-Perry model? In an application to counties of Florida, Smith and Tayman (2003) found that forecasts created from the Hamilton-Perry and cohort-component models were of similar accuracy. However, in an evaluation of various cohort projection models applied to local government areas of New South Wales, Australia, it was found that the Hamilton-Perry model produced age-sex forecasts which were much less accurate than those of the bi-regional cohort-component model (author reference). But when constrained to independent forecasts of total population, the forecast accuracy of the Hamilton-Perry age-sex forecasts improved substantially, coming close to the accuracy of the bi-regional cohort-component model (except at the highest ages). Several other studies have also found that constraining to independent total populations improves the accuracy of small area age-sex population forecasts (e.g., Baker et al. 2020; Reinhold and Thomsen 2015; Tayman et al. 2021).

In addition to constraining to total population forecasts, other implementations of the Hamilton-Perry model have been assessed (author reference), with many of the proposed improvements yielding more accurate forecasts. Some of these consist of modifications to input data, some to forecasts output, and some to model specification. Tayman and Swanson (2017) found that trending small area Cohort Change Ratios and Child/Woman Ratios in line with projected changes in those ratios at the State level improved accuracy. Inoue (2017) reported that smaller forecast errors were obtained by smoothing small area Cohort Change Ratios using a combination of ratios for the small area and those of the larger encompassing region. In addition, Baker et al. (2014) demonstrated that averaging Hamilton-Perry forecasts for urban census tracts with those of neighbouring tracts improved accuracy. In terms of alternative model specification, Hauer (2019) employed a combination of two types of Hamilton-Perry model in county forecasts for the US. Declining populations were projected using Cohort Change Ratios, whilst growing populations were projected using Cohort Change Differences in order to prevent excessively rapid projected growth.

This paper presents findings which build on the recent research noted above by evaluating several implementations of the Hamilton-Perry model to determine which provides the most accurate forecasts. By applying the model to small areas of Australia, the study also offers one of the few examples of the use of this model outside the US. The study was based around four research questions.

1) Are more accurate forecasts obtained from the standard Hamilton-Perry model with Cohort Change Ratios or an alternative version proposed here which uses a mix of Cohort Change Ratios and Cohort Change Differences for each small area?
2) Is it worth smoothing the age profiles of Cohort Change Ratios and Cohort Change Differences?

3) Does constraining to independent small area population totals, and national populations by age and sex, improve accuracy?
4) Do the two Hamilton-Perry models give better forecasts than a No Change model?

The first question concerns the performance of the standard Hamilton-Perry model with that of a modified type. This alternative Hamilton-Perry model prevents rapid growth amongst growing cohorts and excessive decline amongst declining cohorts. It also allows cohorts to grow from an initial population of zero, which is not possible with a Cohort Change Ratio (due to a denominator of zero). It is similar to Hauer's (2019) approach which selects a Cohort Change Ratio or a Cohort Change Difference model depending on whether an area is growing or declining. It also mirrors the composite net migration cohort-component model evaluated by [author reference] where migration rates were used if cohort net migration was negative and net migration numbers if it was positive. This composite model proved more accurate than a cohort-component model using only net migration rates.

The second question deals with age profile smoothing. Inoue (2017) reported gains in accuracy from smoothing age profiles of Cohort Change Ratios, but smoothing has been given little attention in Hamilton-Perry forecasts, so it would be useful to assess the extent to which it affects the forecasts. In this study we assess a simpler smoothing approach than that devised by Inoue.

As mentioned earlier, previous research shows that constraining to independent forecasts of total population can substantially improve accuracy. To answer the third question, we evaluate the effect of constraining to both small area population totals and national forecasts by age and sex. Constraining to national age-sex forecasts gives the small area forecasts the useful property of being fully consistent with national forecasts.

Finally, we are interested in whether the two types of Hamilton-Perry model out-perform a very simple benchmark forecast of no change. For the Hamilton-Perry models to be worth using, they need to perform better than this basic benchmark.

Following this introduction, we describe the forecast methods, input data, forecast variants, and forecast error evaluation metrics. In the Results section we focus on total population forecast errors, errors in the geographical distribution of population across the country, age-sex forecast errors, and errors in forecasting the distribution of the population across age-sex groups within individual areas. The Discussion and Conclusions section includes the key findings of the study, our recommendations for a modified forecasting model, as well as limitations of the study.

2. Data and methods

2.1. Projection models

Two types of Hamilton-Perry cohort projection model were evaluated in this study. The first is a standard version which utilises the Cohort Change Ratio, defined as the ratio of a cohort's population to its size five years earlier when it was five years younger (Baker et al. 2017). This version of the model is labelled **Hamilton-Perry CCR** in this paper. The projected population in age group a+5 at time t+5 is calculated by multiplying the population of the cohort five years earlier, at time t and in age group a, by the Cohort Change Ratio:

$$P^{i}_{s,a+5}(t+5) = P^{i}_{s,a}(t)\ \ CCR^{i}_{s,a\to a+5}(t \to t+5) \quad (1)$$

where

P denotes population

CCR Cohort Change Ratio

t a point in time

$t \to t+5$ the five year forecast interval from time t to t+5

i a small area

s sex

a five year age group (e.g. ages 0-4, 5-9, 10-14)

$a \to a+5$ the cohort aged a at the start of the forecast interval and aged a+5 at the end.

Figure 1 illustrates the Lexis age-time space of equation 1. The thick black vertical lines represent the cohort populations at times t and t+5, while the shaded parallelogram denotes period-cohort population change. The Cohort Change Ratio is a measure of the combined net effect of mortality and migration experienced by the cohort over a period of time. It is typically estimated from population estimates by sex and age group for the jump-off year and five years earlier. For example, the CCR for the cohort of sex s aged 0-4 at time t-5 and aged 5-9 at time t is:

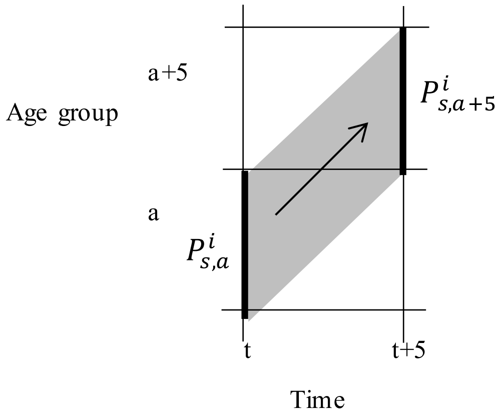

Figure 1: Lexis diagram representation of cohort progression in the Hamilton-Perry model

$$CCR_{s,0-4\to5-9}(t-5\to t)=\frac{P_{s,5-9}(t)}{P_{s,0-4}(t-5)} \qquad (2)$$

Population estimates need to be available up to an age group five years higher than is required in the forecast output. In this study, the available estimates extend by five year age group up to 80-84 and then 85+, so the highest Cohort Change Ratio that can be calculated is for the 80+ to 85+ period-cohort. It is estimated as:

$$CCR_{s,80+\to85+}(t-5\to t)=\frac{P_{s,85+}(t)}{P_{s,80+}(t-5)} \qquad (3)$$

In the projection calculations, the projected end-of-interval cohort populations aged 80-84 (aged 75-79 at the start) and 85+ (aged 80+ at the start) are summed to become the 80+ start-of-interval population for the next forecast interval.

The 0-4 age group is projected differently. Some versions of the Hamilton-Perry model apply the births projection method of the conventional cohort-component model in which age-specific fertility rates are multiplied by age-specific female populations (Hamilton and Perry 1962). The cohort of newly-born infants is then multiplied by an infant CCR to give the projected population aged 0-4. Other versions employ the Child/Woman Ratio (CWR), the ratio of the end-of-interval population aged 0-4 to the female population of childbearing age, e.g. ages 15-49. Here, the population of 0-4 year olds is projected by multiplying the CWR by the female population aged 15-49, which is then multiplied by a sex-specific proportion. The projection equation is:

$$P^i_{s,0-4}(t+5)=P^i_{f,15-49}(t+5)\ \ CWR^i(t+5)\ \ prop^i_s(t+5) \qquad (4)$$

where

f denotes female

CWR Child/Woman Ratio

$prop$ proportion of each sex

The CWR is usually estimated from the jump-off year population.

The alternative version of the Hamilton-Perry model uses a mix of Cohort Change Ratios and Cohort Change Differences (CCDs). If the CCR is less than 1.0, indicating cohort population decline, the cohort population is projected using CCRs. If the CCR is greater than 1.0, indicating cohort growth, it is projected using CCDs. It is referred to as **Hamilton-Perry CCR/CCD** in this paper. This version of the model may be expressed as:

If $CCR^{i}_{s,a\to a+5}(t \to t+5) < 1.0$ then:

$$P^{i}_{s,a+5}(t+5) = P^{i}_{s,a}(t)\ \ CCR^{i}_{s,a\to a+5}(t \to t+5) \qquad (5)$$

Otherwise:

$$P^{i}_{s,a+5}(t+5) = P^{i}_{s,a}(t)\ +\ CCD^{i}_{s,a\to a+5}(t \to t+5)$$

where

CCD denotes Cohort Change Difference.

For many individual small areas, forecasts will be calculated using CCRs for some cohorts and CCDs for others. The 0-4 year old population is projected in the same way as the standard Hamilton-Perry model described above.

A third 'projection model' was included in this study. This is a **No Change** forecast in which the future population by age-sex group is assumed to be the same as in the jump-off year. This forecast acts as a baseline to enable the value of the Hamilton-Perry forecasts to be measured. In order for it to be worth preparing Hamilton-Perry forecasts, they must prove more accurate than a No Change forecast.

In several forecast variants (described below), the age-sex forecasts from the three projection models were constrained to independent small area total populations. These total populations were generated by taking the mean of forecasts generated by the following four models:

1) a Constant Share of Population model, in which each small area's share of the national population is assumed to remain constant from the jump-off year. Forecasts are created by multiplying this share by an independent national population forecast.
2) a composite Linear/Exponential model, in which a linear model is applied if growth over the base period is positive and exponential if it is negative (Rayer 2008). Constraining to the national population total is then applied.
3) a Variable Share of Growth model, where initial small area population growth is generated by the Linear/Exponential model, which is then adjusted to national projected growth using the plus-minus method (author reference).

4) a Modified Exponential model, which comprises an exponential model modified to incorporate floor and ceiling limits to prevent excessive growth or decline (Baker et al. 2008). Constraining to national projected population totals is then applied.

Models 2 – 4 use rates or parameters estimated over the previous ten years, and all models are constrained to an independent national forecast. Recent work has shown that the mean of these four models yields respectable quality short-term small area population forecasts in Australia (Grossman et al. 2021).

2.2. Forecast variants and input data

Population 'forecasts' for 2006-16 were prepared for small SA2 areas of Australia (defined by the 2011 Australian Statistical Geography Standard). SA2 areas were designed by the Australian Bureau of Statistics to contain populations generally within the range 3,000 to 30,000, though some populations fall outside these limits. In 2006 the median population of an SA2 area was 8,426. Those SA2 areas with fewer than 1,000 residents in 2006 were judged too small to give reliable cohort change ratios or differences, and were aggregated into one 'remainder' area for each State and Territory. A total of 2,054 geographical areas was included in our study. Population estimates data by age and sex for a consistent set of geographical areas was only available for the years 2001 to 2016, limiting us to a base period for calculating CCRs and CCDs of 2001-06 and a forecast horizon of 2006-16.

Ten forecast variants were produced, which are summarised in Table 1. Forecast variants were prepared to evaluate the effects of:

(a) using Hamilton-Perry CCR versus Hamilton-Perry CCR/CCD,

(b) constraining to independent population forecasts,

(c) smoothing the age profiles of CCRs and CCDs, and

(d) using Hamilton-Perry models compared to a No Change assumption.

The first set of forecasts (labelled 1a to 1e in Table 1) was not subject to any constraining. All values of CCRs, CCDs and CWRs were assumed to remain constant from the base period. For the second set (2a to 2e), forecasts were initially calculated in the same way, but were adjusted using iterative proportional fitting to be consistent with (i) national forecasts by sex and age group and (ii) SA2 projected total populations. Figure 2 illustrates the population matrix layout used in constraining, with initial forecasts depicted by the lightly shaded cells and constraints shown by the dark shaded cells. The national population constraints are represented by the row totals on the right-hand side, while the SA2 total population constraints are depicted by the column totals at the bottom of the diagram. The national forecasts were the ABS 2006-based Series B (medium series) forecasts (ABS 2008), while the SA2 total population forecasts were created with the model averaging approach (themselves constrained to the ABS national forecast totals).

Table 1: Population forecast variants

	Model type				
	Hamilton-Perry CCR		**Hamilton-Perry CCR/CCD**		**No Change**
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
1. Unconstrained	1a	1b	1c	1d	1e
2. Constrained to national forecasts & SA2 total population forecasts	2a	2b	2c	2d	2e

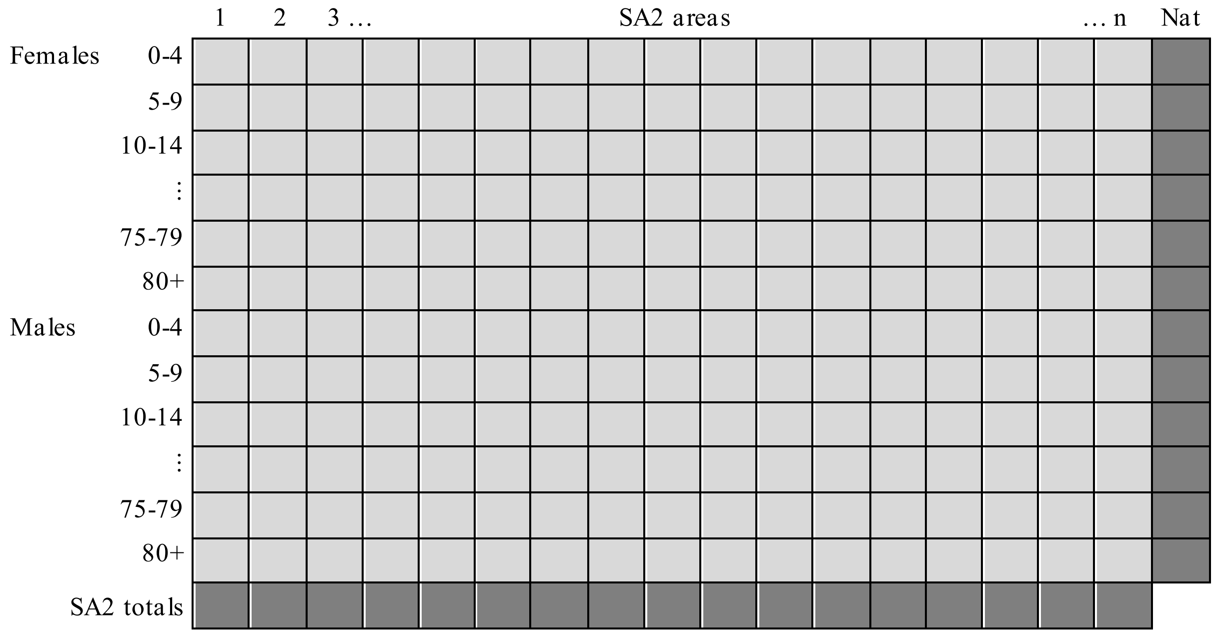

Figure 2: The layout of the SA2 age-sex population matrix used in the constrained forecasts
Note: Nat = national forecasts. Light shading indicates initial forecasts; dark shading depicts constraining populations.

The smoothing applied to the CCR age profiles consisted of an adaptation of de Beer's TOPALS approach using rate ratios (de Beer 2012). In essence this approach involves (i) taking a model set of age-specific rates (e.g., national rates), (ii) calculating age-specific local/model rate ratios, (iii) smoothing the rate ratios over age, and then (iv) creating smoothed local rates by multiplying the model rates by the smoothed rate ratios. For each SA2 area, the ratio of the area's CCR to the national CCR was calculated by sex and period-cohort[1]. For period-cohorts 0-4 – 5-9 to 40-44 – 45-49 the mean of male and female CCR rate ratios was assumed for both sexes. Differences in mortality and net migration between males and females should, for most areas, be fairly small

[1] A period-cohort refers to the Lexis space defined by a cohort over a specified period (Figure 1). We use notation such as 0-4 – 5-9, which refers to the cohort aged 0-4 at the start of the projection interval and aged 5-9 at the end.

across the childhood and early to mid-adult ages. Smoothing was not applied across age in order to allow area-specific net migration patterns to remain. At higher ages, CCR ratios were smoothed over age for each sex because migration patterns often change more gradually over age at older ages. For period-cohorts 45-49 – 50-54 and above, linear splines were applied to sex-specific CCR ratios, with knots selected at every other period-cohort using a three-term average (the average of the period-cohort in question and the ones immediately younger and older). For the final 80+ – 85+ period-cohort the average of the final four period-cohort CCR ratios was used due to the volatility of the data at these ages. Smoothed CCRs were then calculated as national CCRs multiplied by the smoothed CCR ratios. For the age profile of CCDs, CCR local/national ratios could not be used and smoothing was applied directly to the CCD numbers. The average CCD of males and females was assumed for each sex up to the 40-44 – 45-49 period-cohort, and linear splines were applied to sex-specific CCDs for period-cohorts 45-49 – 50-54 upwards. A worked example of the smoothing procedure is provided in the supplementary materials accompanying this paper.

Input data for the SA2 forecasts consisted of Estimated Resident Populations (ERPs) for 30th June 2001 and 2006 by sex and five year age group from 0-4 to 80-84 and 85+. These populations were used to calculate CCR and CCDs, with the 2006 populations also forming the jump-off populations for the forecasts. Equivalent SA2 and national ERPs for 2011 and 2016 were also obtained for use in constraining and to assess forecast errors. Total SA2 ERPs for 1996 and 2006 were downloaded to estimate the rates and parameters for the four models used to create SA2 total projected populations. All this data is freely available and was obtained from the ABS.Stat website, https://stat.data.abs.gov.au/. The ABS 2006-based population forecasts data (ABS 2008) is also freely available from the ABS website (see https://tinyurl.com/wf8j9ext). We have collected all the data needed to create the forecasts, and then evaluate them, in one file. It is available at https://doi.org/10.6084/m9.figshare.16822495.v1.

2.4. Error measures

The SA2 area population 'forecasts' were evaluated by calculating Percentage Error (PE), defined as:

$$PE = \frac{(F - ERP)}{ERP} \, 100$$

where

F denotes the forecast

ERP Estimated Resident Population.

Errors are summarised using Median Absolute Percentage Error (MedAPE), the middle of the distribution of Absolute Percentage Errors, and Mean Absolute Percentage Error (MAPE). MedAPE is our preferred measure of 'typical' error because it is not influenced the large outliers; we report *mean* error values in appendix tables to allow comparisons with studies which present mean errors. Bias – indicating whether forecasts are too high or too low overall – is summarised using Median Percentage Error, the middle of the distribution of signed Percentage Error values. In appendix tables we also report Mean Percentage Error.

The Index of Dissimilarity (IoD) is used to quantify the error in forecasting population distributions. We include this measure because some uses of population forecasts depend on the projected *distribution* of population rather than population *numbers*. The IoD is defined as (Rowland 2003):

$$IoD = ½ \sum_i |a_i - b_i|$$

where

a_i is the share of observation i amongst a set of values A

b_i is the share of observation i amongst a set of values B.

We apply it to evaluate the difference between forecast and actual age-sex distributions of individual SA2 populations, and also the error in forecasting the geographical distribution of SA2 total populations. For the latter example, a_i would represent the projected share of the national population in SA2 i while b_i would be the actual share of the national population in SA2 i. Values of the index range between 0 and 1, with values close to 0 indicating low errors in forecasting the actual distribution. The IoD can be interpreted as the share of the forecast population that would need to move between categories to achieve the actual distribution, e.g., a value of 0.05 indicates that 5% of the population would need to move between categories to achieve an equal distribution.

3. Results

3.1. Total population forecasts

The top panel of Table 2 presents Median Absolute Percentage Errors (MedAPEs) for total SA2 populations at forecast horizons of 5 and 10 years. At 5 years out there is little difference between the unconstrained Hamilton-Perry forecasts (top row of the table), with MedAPEs varying only between 4.7% for the Hamilton-Perry CCR model and 4.6% for the CCR/CCD version. The unconstrained No Change forecast, not surprisingly, produces higher errors because Australia's small area population totals are far from static. The constrained forecasts (series 2) are more accurate, and all take the same error values because of constraining to the same set of averaged model forecasts.

After 10 years, there is a little more differentiation in median errors between the CCR and CCR/CCD versions of the Hamilton-Perry model in the unconstrained forecasts, with a small gain in accuracy achieved by using the CCR/CCD version. Again, the constrained forecast errors are substantially more accurate than all the unconstrained forecasts. The MedAPE of 6.3% for small area population totals after 10 years is a respectable level of error and comparable to that found in other forecasting evaluation studies using SA2 areas (e.g., author reference). For both 5 and 10 year forecast horizons, there is no effect from whether the CCR and CCD age profiles are smoothed or not. However, smoothing would be expected to have more impact on age-sex population forecasts rather than population totals.

Absolute Percentage Errors at the 90th percentile of the error distribution are shown in the second panel of Table 2. They are, not surprisingly, much higher than the median errors, but they follow the same pattern: in the unconstrained forecasts the 90th percentile APEs are a little lower in the Hamilton-Perry CCR/CCD model, and lower for all the constrained forecasts.

Table 2 also presents median percentage errors which indicate the bias of the forecasts. Not surprisingly, the unconstrained No Change forecasts were biased downwards (forecasting populations which, overall, were too low) because they failed to include population growth in the Australian population. The unconstrained Hamilton-Perry forecasts (1a to 1d) were slightly too low. However, the constrained (series 2) forecasts experienced little bias, even though the 2006-based national population forecast, to which they were constrained, under-forecast the total Australian population by 1.9% by 2016.

The bottom panel of Table 2 indicates the percentage of SA2 areas from the unconstrained forecasts with total population APEs which were lower than those of the total populations used in the constrained forecasts (2a to 2e). For all unconstrained forecasts, only a minority of areas achieved lower APEs than the constrained forecasts, confirming that constraining to independent total populations is worthwhile overall. (Values are not shown for the constrained set of forecasts because all total populations are the same.)

Table 2: Error and bias of SA2 total population forecasts

	Model type				
	Hamilton-Perry CCR		**Hamilton-Perry CCR/CCD**		**No Change**
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
MedAPE (%)					
5 year horizon					
1. Unconstrained	4.7	4.7	4.6	4.6	5.8
2. Constrained to national forecasts & SA2 total population forecasts	3.3	3.3	3.3	3.3	3.3
10 year horizon					
1. Unconstrained	8.9	8.9	8.4	8.5	10.0
2. Constrained to national forecasts & SA2 total population forecasts	6.3	6.3	6.3	6.3	6.3
90th percentile APE (%)					
5 year horizon					
1. Unconstrained	13.6	13.5	12.3	12.3	16.2
2. Constrained to national forecasts & SA2 total population forecasts	11.9	11.9	11.9	11.9	11.9
10 year horizon					
1. Unconstrained	26.2	26.1	22.1	22.0	28.3
2. Constrained to national forecasts & SA2 total population forecasts	20.6	20.6	20.6	20.6	20.6
Median PE (%)					
5 year horizon					
1. Unconstrained	-2.5	-2.5	-2.9	-2.9	-5.5
2. Constrained to national forecasts & SA2 total population forecasts	-0.2	-0.2	-0.2	-0.2	-0.2
10 year horizon					
1. Unconstrained	-4.4	-4.4	-5.2	-5.3	-9.5
2. Constrained to national forecasts & SA2 total population forecasts	0.7	0.7	0.7	0.7	0.7
% of SA2 areas with lower APE than constrained forecasts					
5 year horizon					
1. Unconstrained	38.6	38.3	39.8	39.8	27.7
10 year horizon					
1. Unconstrained	37.0	36.9	40.0	39.8	31.0

Notes. The lowest errors are indicated by shading. Error values are the same for all constrained forecasts due to SA2 total population constraining. Constrained forecasts are therefore not shown in the bottom panel because the total populations are the same across all constrained forecast series.

Table 3 reports errors in forecasting the geographical distribution of SA2 total populations across Australia as measured by the Index of Dissimilarity. This ignores the forecast error of projected populations, focusing instead on the error in forecasting the *proportional distribution* of the national population across small areas. At a 5 year forecast horizon, the Hamilton-Perry CCR/CCD variants (smoothed and unsmoothed; constrained and unconstrained) all produce forecasts of the geographical distribution of SA2 populations that are very similar and the most accurate.

After 10 years, the difference between the two unconstrained Hamilton-Perry types was considerable (1a and 1b versus 1c and 1d). Interestingly, the unconstrained No Change models (1e and 2e) produced forecast distributions which lie between the CCR and CCR/CCD versions in terms of accuracy at both 5 and 10 year horizons. Constraining markedly improved the accuracy of the Hamilton-Perry CCR version, but made little difference to the forecasts generated by the CCR/CCD type, which were already quite good in the unconstrained forecasts.

Table 3: Errors in forecasting the geographical distribution of SA2 area total populations as measured by the Index of Dissimilarity

	Model type				
	Hamilton-Perry CCR		**Hamilton-Perry CCR/CCD**		**No Change**
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
5 year horizon					
1. Unconstrained	0.0399	0.0399	0.0272	0.0274	0.0338
2. Constrained to national forecasts & SA2 total population forecasts	0.0275	0.0275	0.0275	0.0275	0.0275
10 year horizon					
1. Unconstrained	0.1264	0.1256	0.0531	0.0534	0.0657
2. Constrained to national forecasts & SA2 total population forecasts	0.0527	0.0527	0.0527	0.0527	0.0527

Note: Index of Dissimilarity values are the same for all constrained forecasts due to SA2 total population constraining. The lowest errors are indicated by shading.

3.2. Age-sex population forecasts

The median of absolute percentage errors across all SA2 areas and all age-sex groups are reported in the top panel of Table 4. At both 5 and 10 year forecast horizons, the Hamilton-Perry CCR/CCD model produced errors a little lower than those of the CCR version (by 0.4-1.0 percentage points after 10 years). For example, after 10 years the MedAPE for the unconstrained CCR model without CCR age profile smoothing (1a) is 14.2%, while for the equivalent CCR/CCD version (1c) is 13.3%. There is a similar lowering of error in using smoothed CCR or CCD age profiles compared to unsmoothed age profiles. For example, the unconstrained CCR model with smoothed age profiles (1b) produced a MedAPE of 13.4% after 10 years (compared to 14.2% without smoothing). Constraining to independent population forecasts (series 2) also reduced errors modestly in all cases (by 0.8 to 1.3 percentage points after 10 years). The lowest median errors of all the forecast variants were obtained by the Hamilton-Perry CCR/CCD model constrained to independent forecasts using smoothed CCR and CCD age profiles (2d).

Errors of the No Change forecasts were highest, indicating that forecasting age structure change is generally sensible. The better-performing constrained No Change forecast (2e) experienced a higher MedAPE than all the Hamilton-Perry variant forecasts (but only by a small margin after 10 years).

At the 90th percentile of the APE distribution (shown in the second panel of Table 4), age-sex forecast errors are approximately three times the median values. Errors are a little lower in the forecast variants with smoothed CCR and CCD age profiles, and with the CCR/CCD model than the CCR version. However, constraining shows mixed results: a small improvement in accuracy occurs by constraining the CCR model forecasts, but not those of the CCR/CCD model. Interestingly, after 10 years, the unconstrained Hamilton-Perry CCR model (1a) generated a slightly higher 90th percentile APE than the constrained No Change model (2e).

The third panel of Table 4 reports forecast bias, with the values shown referring to the median of percentage errors across all SA2 areas and all age-sex groups. In the unconstrained forecasts, the age-sex populations are slightly under-forecast, overall. In the unconstrained No Change forecast (1e) the under-forecasting is especially severe at the older ages (not shown) because the model does not project the growth of older populations through cohort progression. The table shows that the application of constraining (series 2 forecasts) removes most of the bias; it also substantially reduces the variation in bias between age-sex groups (not shown).

Table 4: Error and bias of SA2 age-sex population forecasts

	Model type				
	Hamilton-Perry CCR		**Hamilton-Perry CCR/CCD**		**No Change**
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
MedAPE (%)					
5 year horizon					
1. Unconstrained	9.5	8.8	9.0	8.5	12.7
2. Constrained to national forecasts & SA2 total population forecasts	9.0	8.3	8.8	8.2	10.4
10 year horizon					
1. Unconstrained	14.2	13.4	13.3	12.6	18.6
2. Constrained to national forecasts & SA2 total population forecasts	13.0	12.1	12.5	11.7	14.3
90th percentile APE (%)					
5 year horizon					
1. Unconstrained	28.7	26.7	26.3	24.6	33.3
2. Constrained to national forecasts & SA2 total population forecasts	27.7	25.7	26.7	24.9	29.6
10 year horizon					
1. Unconstrained	42.2	40.2	36.6	35.0	46.9
2. Constrained to national forecasts & SA2 total population forecasts	39.5	37.2	37.2	35.4	40.5
Median PE (%)					
5 year horizon					
1. Unconstrained	-2.2	-2.0	-2.6	-2.4	-7.1
2. Constrained to national forecasts & SA2 total population forecasts	-0.6	-0.3	-0.3	-0.2	-0.7
10 year horizon					
1. Unconstrained	-4.6	-4.1	-5.2	-4.9	-12.2
2. Constrained to national forecasts & SA2 total population forecasts	-0.4	0.0	-0.1	0.2	-0.4
% of SA2 areas with lower MedAPE over age and sex than constrained No Change (2e)					
5 year horizon					
1. Unconstrained	60.7	67.0	65.3	71.2	25.5
2. Constrained to national forecasts & SA2 total population forecasts	65.8	74.3	71.2	77.6	n/a
10 year horizon					
1. Unconstrained	49.9	55.7	56.4	61.8	22.9
2. Constrained to national forecasts & SA2 total population forecasts	58.7	65.7	64.5	71.0	n/a

Note: The lowest errors or best results are indicated by shading.

The bottom panel of Table 4 indicates the percentage of SA2 areas where the median APE over age and sex was lower than the equivalent error from the constrained No Change model (2e). Again, the CCR/CCD model with smoothing performs best, achieving a lower age-sex MedAPE than the constrained No Change model in 77.6% of areas after 5 years, and in 71.0% of areas after 10 years.

To illustrate the variation in age-sex-specific forecast errors, Figure 3 shows MedAPEs by age group for females for selected variants after 10 years. The graph includes the worst and best performing of the Hamilton-Perry model variants (1a and 2d) and the two No Change forecasts (1e and 2e). The pattern for males is similar and not shown. Both Hamilton-Perry forecasts exhibit similar age patterns of errors, with a peak in the young adult ages where migration rates are highest and most variable, and higher errors at older ages where population sizes tend to be smaller. The constrained Hamilton-Perry CCR/CCD model with smoothed CCR/CCD age profiles (2d) gave the lowest errors at most ages by a modest margin. The difference between 1a and 2d shows the combined effects of using the CCR/CCD model, smoothing age profiles, and constraining to independent forecasts. There is a positive impact at all ages, though the effect is greatest in the young childhood ages (about 5 percentage points) and smallest at the older ages (1.5 to 3.0 percentage points).

As Figure 3 shows, the patterns of errors are quite different for the No Change variants (1e and 2e). The unconstrained No Change model (1e) generated very high errors at the older ages because this model fails to incorporate population ageing. Although constraining to national age-sex forecasts reduced these errors considerably (2e), errors at most ages above 60 remain higher than those of the Hamilton-Perry variants. However, at many ages below 60 the constrained No Change model produced lower median errors than unconstrained Hamilton-Perry CCR (1a). It also gave slightly lower errors than the constrained Hamilton-Perry CCR/CCD model (2d) at ages 25-29.

Figure 3: MedAPEs by age group for females after 10 years, selected forecast variants
Note: Forecast variants are defined in Table 1.

Errors in forecasting the *distribution* of population across age-sex groups within SA2 areas are summarised by the Index of Dissimilarity measures in Table 5. The Index of Dissimilarity shown here indicates whether the projected SA2 populations are distributed in the correct proportions across age-sex groups (not whether the forecast age-sex population numbers themselves are accurate). In terms of the median of the Index of Dissimilarity distribution (top panel of the table), slightly lower errors were achieved by smoothing CCR and CCD age profiles and by using the CCR/CCD model rather than the CCR version. Constraining either marginally lowered errors or had no impact. A similar effect is observed at the 90th percentile of the distribution (middle panel). Relative to the better-performing No Change forecast (2e), all Hamilton-Perry age-sex distribution forecasts have lower median errors, while most (but not all) errors at the 90th percentile of the Index of Dissimilarity distribution are lower.

Table 5: Errors in forecasting the population age-sex distribution of SA2 areas as measured by the Index of Dissimilarity

	Model type				
	Hamilton-Perry CCR		**Hamilton-Perry CCR/CCD**		**No Change**
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
Median IoD					
5 year horizon					
1. Unconstrained	0.048	0.043	0.046	0.042	0.061
2. Constrained to national forecasts & SA2 total population forecasts	0.047	0.043	0.045	0.041	0.054
10 year horizon					
1. Unconstrained	0.066	0.061	0.061	0.056	0.083
2. Constrained to national forecasts & SA2 total population forecasts	0.065	0.060	0.059	0.055	0.070
90th percentile IoD					
5 year horizon					
1. Unconstrained	0.079	0.073	0.074	0.067	0.087
2. Constrained to national forecasts & SA2 total population forecasts	0.079	0.072	0.073	0.067	0.080
10 year horizon					
1. Unconstrained	0.112	0.102	0.100	0.093	0.122
2. Constrained to national forecasts & SA2 total population forecasts	0.111	0.102	0.097	0.090	0.106
% of SA2 areas with lower IoD than constrained No Change (2e)					
5 year horizon					
1. Unconstrained	65.4	77.8	73.1	82.6	21.6
2. Constrained to national forecasts & SA2 total population forecasts	67.9	80.1	75.9	86.2	n/a
10 year horizon					
1. Unconstrained	53.2	64.5	63.5	72.8	24.3
2. Constrained to national forecasts & SA2 total population forecasts	55.1	67.1	66.2	76.3	n/a

Note: The best values are indicated by shading.

The bottom panel of Table 5 reports the percentage of SA2 areas where the age-sex distribution of an area's population was forecast more accurately according to the Index of Dissimilarity than the better-performing No Change model (2e). At both 5 and 10 year horizons, all Hamilton-Perry variants projected better age-sex distributions than 2e for a majority of areas. The effects of smoothing CCR and CCR age profiles, using the CCR/CCN model, and constraining, all substantially increased the percentage of areas with more accurate projected age-sex distributions.

Table 6 examines how the forecast error in the age-sex distribution varied by population size and base period growth rate for the Model 2d forecasts, which performed best. The patterns by population size and growth rate were very similar for all other Hamilton-Perry variants. The top half of the table shows median Index of Dissimilarity values for each SA2 total population size category. An inverse relationship with population size, found in many forecast error studies (e.g., Tayman 2011) is evident: higher errors can be seen for the smallest populations and lowest errors for the largest. The relationship is non-linear, with errors falling quickly from the smallest population size category to the next, but with smaller declines as population size increases. A reduction in the amount of noise in age-sex CCR and CCD age profiles as population size increases is likely to be at least part of the explanation. The bottom half of the table indicates how median IoD values varied by base period (2001-06) population growth rate. In common with other studies (Tayman 2011), a u-shaped relationship – albeit quite shallow – is evident, with the most accurate age-sex distributions forecast for areas with base period growth rates between 0 and 1% per annum. Areas experiencing decline and high growth rates were slightly less well forecast on average.

Table 6: Errors in forecasting the population age-sex distribution of SA2 areas by population size and growth rate categories as measured by the median Index of Dissimilarity, Model 2d

Jump-off year population size							
	<3,000	3,000-4,999	5,000-6,749	7,500-9,999	10,000-14,999	15,000-19,999	20,000+
5 year horizon	0.0728	0.0554	0.0450	0.0396	0.0347	0.0297	0.0272
10 year horizon	0.0851	0.0710	0.0590	0.0513	0.0467	0.0421	0.0397
Base period annual average growth rate (%)							
	<-1	-1 to 0	0 to 1	1-2	2-3	3-4	4+
5 year horizon	0.0494	0.0391	0.0384	0.0413	0.0430	0.0427	0.0460
10 year horizon	0.0674	0.0545	0.0511	0.0560	0.0573	0.0569	0.0573

Some hints about the origins of poorly forecast age-sex structures were gleaned by mapping Index of Dissimilarity values. Figure 4 shows the geographical variation in Index of Dissimilarity values in 2016 for the age-sex distribution forecast by the best-performing model (2d). Red shading indicates values above the median (less accurate), while green shading represents below-median values (more accurate). While the red shading appears to dominate the map, many of small areas where the age-sex distributions were forecast less well are geographically large 'small' areas in remote Australia. Many of these areas also have smaller than average SA2 populations where the CCRs and CCDs are subject to greater amounts of noise. They also include mining areas where population size can change rapidly with economic conditions and mine openings/closures. Some of the apparent poor forecasting of age-sex structure in remote regions may also be due to the difficulties in enumerating remote and Indigenous communities in the census (ABS 2018). In the major cities, age-sex structure was poorly forecast in rapidly growing residential areas where the age structure of the population undergoes substantial transformation, and in inner-city areas where it remained fairly stable. Values of the IoD for the age-sex distribution forecast are in the fourth quartile for the majority of areas in the Australian Capital Territory where SA2 populations in 2006 were generally smaller than average.

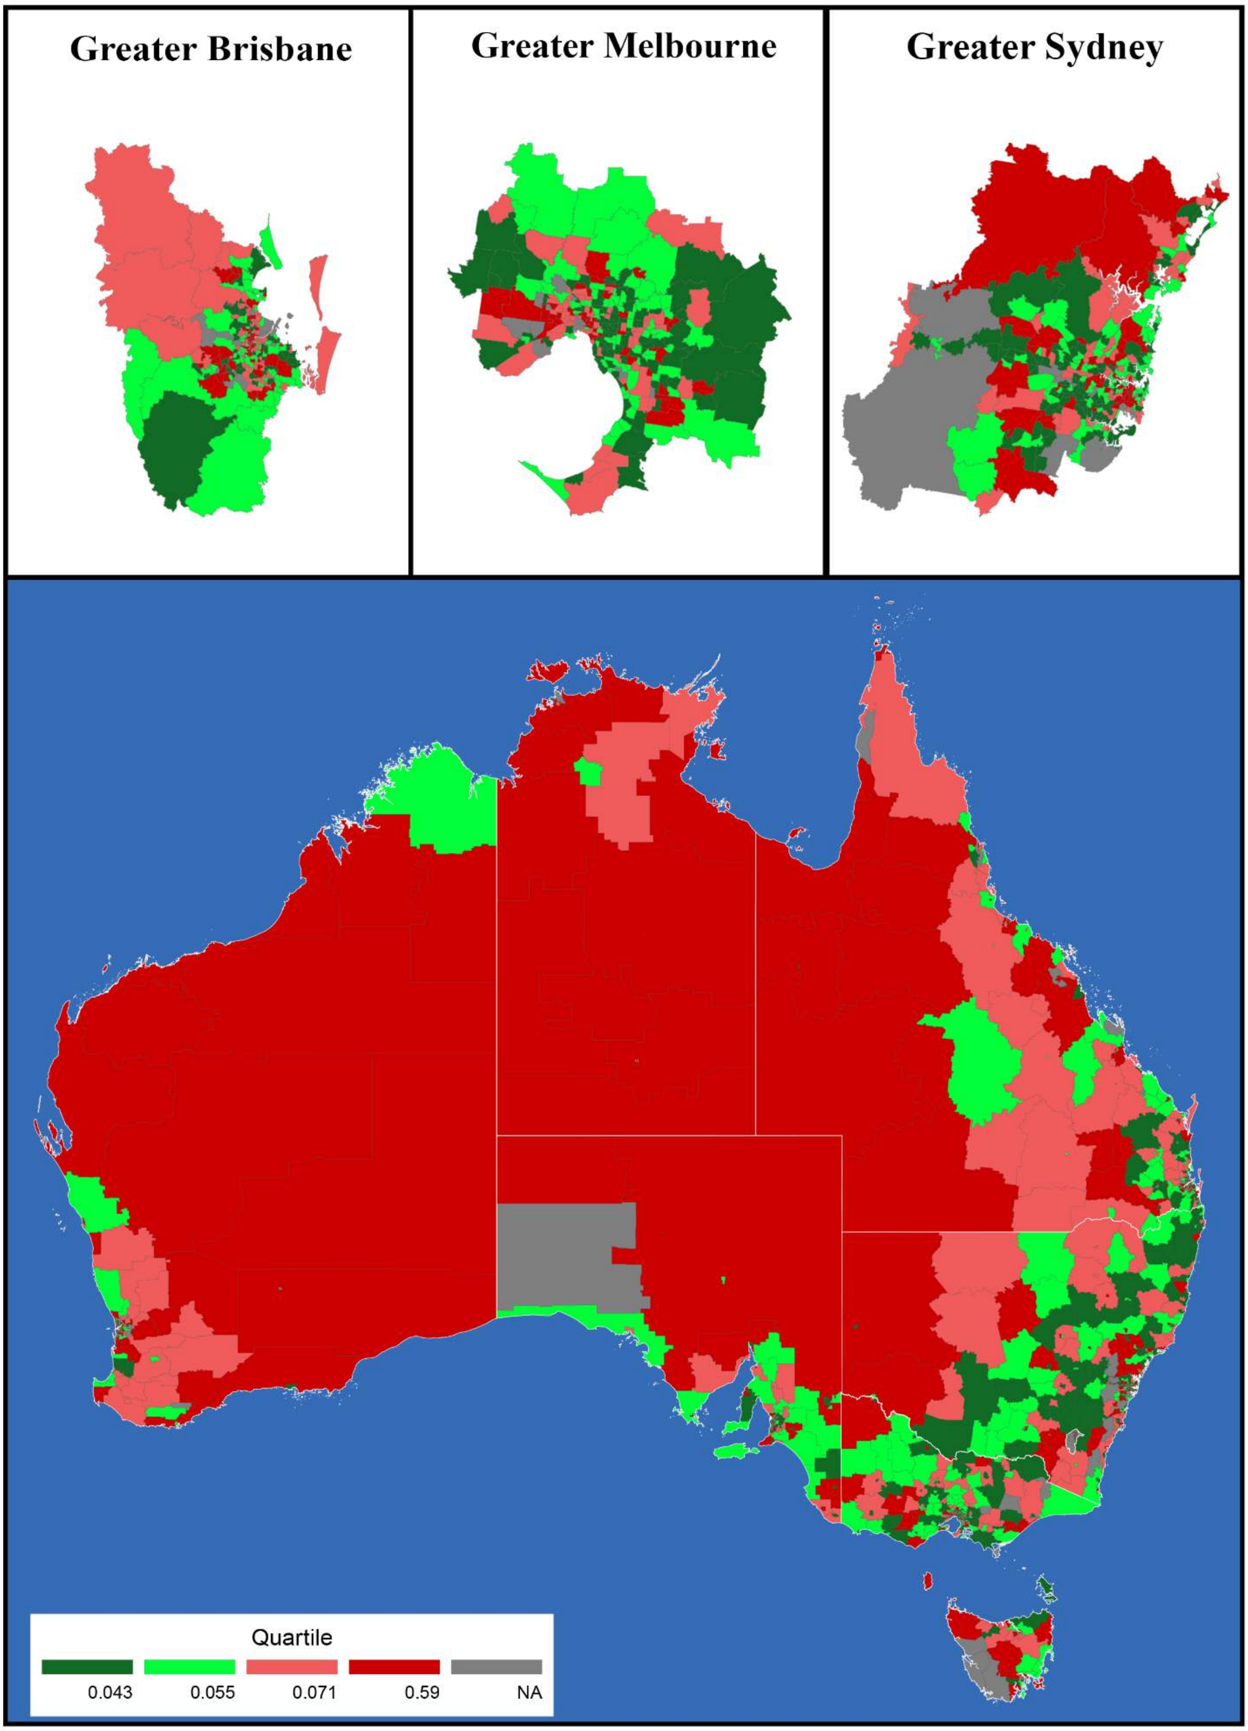

Figure 4: Spatial distribution of Index of Dissimilarity values of the age-sex distribution for model 2d in 2016

To explore the relationship between errors in forecasting the age-sex distribution and potential explanatory variables further, we performed a simple OLS regression with the 2016 age-sex Index of Dissimilarity from Model 2d as the dependent variable. The independent variables were:
(i) the Index of Dissimilarity indicating the difference between the 2001 and 2006 age-sex distributions for each SA2 area;
(ii) the SA2 jump-off population size (logged);
(iii) the percentage of people in each SA2 who lived at a different address 5 years ago; and
(iv) the percentage of the population identifying as Indigenous;
(v) the total population absolute percentage error for 2016;
(vi) the percentage of people living in communal accommodation (e.g., military barracks, prisons, boarding schools); and
(vii) the annual average population growth rate in the 2001-06 base period.
The first two variables act as proxy indicators of the amount of noise in the data (the smaller the population, the noisier the data, and the more variable the age-sex distributions 5 years apart). The third variable measures the extent to which the population is affected by migration. Because migration trends can be quite volatile over time, the more migration experienced by a population, the greater the potential for forecast error. The share of the population Identifying as Indigenous is related to challenges in enumerating the population accurately. The fifth variable, the total population Absolute Percentage Error in 2016, refers to the influence of total forecast population change being too high or too low on age distribution errors. The sixth variable is the share of the SA2 population living in communal accommodation. Populations in these dwellings tend to experience high inward and outward mobility and change little in age-sex structure over time. The final variable is the base period population growth rate, which other studies have found to be associated with error (e.g., Tayman et al. 2011). Variables describing the percentage of the workforce in mining and construction employment, the percentage of higher education students, and the remoteness classification of each SA2 were found to add little to the model and were excluded.

Table 7 summarises the regression results. Overall, the results of the regression are statistically significant (R^2=0.524, $F(6, 2047) = 376.2$, $p < 2.2e-16$). The model accounts for just over half the variation in the age-sex IoD. The right-hand column of the table indicates the percentage contribution to R-squared calculated using Shapley's decomposition (Shapley 1953), indicating that the 2001-06 age-sex IOD, the natural log of the jump-off population, and the rate of mobility (% different address) contribute the most. But about half the variation remains unexplained. As is often observed, small areas can experience major shifts in their demographic trajectories over time, and therefore change in ways which differ substantially from that recorded during the base period used in population forecasting. For comparative purposes, we also present an OLS regression where MedAPE over age and sex is the dependent variable (Appendix Table A5).

Table 7: Regression results for age-sex distribution Index of Dissimilarity, Model 2d, 2016

Variable	Coefficient	Standard error	p value	% contribution†
Constant	1.067e-01	9.599e-03	< 2e-16	
2001-06 age-sex IoD	4.867e-01	2.898e-02	< 2e-16	30.3
Ln(jump-off population)	-1.244e-02	9.016e-04	< 2e-16	25.7
% different address	1.659e-03	9.365e-05	< 2e-16	17.0
% Indigenous pop	1.323e-04	4.737e-05	0.00527	2.9
Total pop APE	6.390e-04	4.627e-05	< 2e-16	13.4
% communal dwellings	9.221e-04	1.173e-04	6.08e-15	9.0
Base period growth rate	-1.573e-03	1.790e-04	< 2e-16	1.7
No. of observations	2,054			
Adjusted R-squared	0.540			

Notes. † Shapley decomposition of the contribution to R-squared

4. Discussion & Conclusions

4.1. Summary of key findings

This study evaluated the standard version of the Hamilton-Perry projection model for small areas of Australia (Hamilton-Perry CCR) and an alternative version with a mix of CCRs and CCDs (Hamilton-Perry CCR/CCD). It included forecast variants which were constrained to independent small area total population and national age-sex forecasts, and where CCR and CCD age profiles had been smoothed. Here, we summarise answers to each of our research questions.

Are more accurate forecasts obtained from the standard Hamilton-Perry CCR model or the alternative CCR/CCD version?

In this study, more accurate forecasts were produced by the alternative Hamilton-Perry CCR/CCD model. In the unconstrained set of forecasts, errors in projecting total populations with the CCR/CCD version were slightly lower at both the median and 90th percentile of the APE distribution (Table 2). The geographical distribution of small area total populations across the country was considerably more accurately forecast by the CCR/CCD model in the unconstrained forecasts (Table 3). (Errors for total populations were identical for all variants in the constrained set of forecasts due to constraining to averaged model SA2 population totals.) Overall, age-sex forecasts were slightly more accurately forecast with the CCR/CCD model, for both unconstrained and constrained forecast variants. This was the case both for age-sex APEs (Table 4) and the age-sex distribution of the population as measured by the Index of Dissimilarity (Table 5).

Is it worth smoothing the age profiles of CCRs and CCDs?

Yes. The simple smoothing method we applied to CCRs and CCDs produced modest gains in the accuracy of age-sex population forecasts. This was demonstrated in terms of both MedAPEs and 90th percentile APEs (Table 4), and the distribution of population across age-sex groups as measured by the Index of Dissimilarity (Table 5). For all categories of SA2 population size, the smoothed age profile CCR/CCD model proved best.

Does constraining to independent small area population totals, and national populations by age and sex, improve accuracy?

Overall, yes. Constraining considerably improved the accuracy of SA2 total populations from all forecast variants (Table 2). In terms of forecasts of the geographical distribution of population totals, constraining substantially improved the Hamilton-Perry CCR forecasts, but had almost no impact on the CCR/CCD model's performance (Table 3) which was already quite good. Constraining modestly improved the accuracy of age-sex forecasts as measured by MedAPE (Table 4) and the accuracy of the age-sex distribution of the population as measured by the Index of Dissimilarity (Table 5).

Do the two Hamilton-Perry models give better forecasts than a No Change model?

Mostly. In terms of total populations, the unconstrained Hamilton-Perry CCR and CCR/CCD models (1a to 1d) produced lower MedAPEs than the unconstrained No Change model (1e) (Table 2). (The constrained versions all gave the same MedAPE due to SA2 total population constraining.) For the age-sex forecasts, all eight Hamilton-Perry variants (1a to 1d; 2a to 2d) achieved lower age-sex MedAPEs than both constrained and unconstrained No Change models (but only by a tiny margin for variant 1a after 10 years) (Table 4). The Hamilton-Perry models also gave lower median Index of Dissimilarity values for the age-sex distribution forecast than the No Change models (Table 5). However, these findings refer to error across age-sex groups overall. The picture varies considerably when examining individual age groups (Figure 3). The reduction in error by switching from the No Change models to the Hamilton-Perry models was greatest at the older ages because of the inclusion of population ageing when cohort progression is modelled. But below age 50, the difference in forecast performance between these two sets of models was smaller. At most age groups below age 50, the worst-performing Hamilton-Perry variant (1a) experienced slightly higher age-specific MedAPEs than the better-performing constrained No Change model (2e). However, the best-performing Hamilton-Perry variant (2d) achieved lower age-specific MedAPEs than the constrained No Change model (2e) at every age except 25-29; this was the case for both females (Figure 3) and males (not shown).

4.2. Limitations

The limitations of this study include the evaluation of the chosen forecast variants over a 10 year horizon only, and for a single 10 year period only, due to data availability. It is important to emphasise that the characteristics of the forecasts reported here for the 2006-16 period may not be replicated in other periods. The evaluation also measured 'forecasts' relative to Estimated Resident Populations on the assumption that the latter were precise measures of the true resident population. In reality, they are estimates, not precise counts. In addition, our evaluation applies to Australia only, and findings are likely to differ to some extent in other countries.

4.3 Suggested forecasting approach

Based on the findings reported here, our suggestion for those tasked with creating forecasts is that, where it is not possible to implement a full cohort-component model, the Hamilton-Perry CCR/CCD model with smoothed CCR and CCD age profiles and constrained to independent forecasts (2d) represents a reasonable option for short-term small area forecasts up to 10 years ahead. However, it is important to be aware of the model's limitations. These include: large forecast errors for some areas, especially those with the smallest populations and those areas experiencing a radical change in demographic growth and patterns; poor results from averaging male and female CCRs or CCDs for some areas; the lack of projected demographic components of change; and uncertain long-term forecast performance. We recommend that the following practical modifications be considered to overcome some of these limitations.

First, age-sex structure modifications are recommended for certain areas. While the population age-sex structure is reasonably forecast for many small areas, there are some areas where this will not be the case. For some areas, the population age structure changes little over time because the dwelling mix attracts a mobile population which moves in and out at high rates but is fairly static in its demographic characteristics. These areas include inner city areas dominated by higher education students and young professionals, boarding schools, defence facilities, prisons, and large communal accommodation facilities. Sometimes base period CCRs and CCDs fail to fully capture these trends. For selected areas, therefore, the population at some ages could be calculated using the No Change model. This will maintain a steady age structure. For example, the SA2 area of Adelaide, which covers the central business district of Adelaide, South Australia, is dominated by a young adult population. Figure 4 below illustrates its forecast in 2016 based on 2001-06 CCRs and CCDs using model 2d together with a modified forecast where the population under age 65 (pre-constraining) is held constant from the jump-off year. The modified forecast is far from perfect, but achieves a projected age-sex structure much closer to the recorded Estimated Resident Population in 2016 than the model 2d forecast.

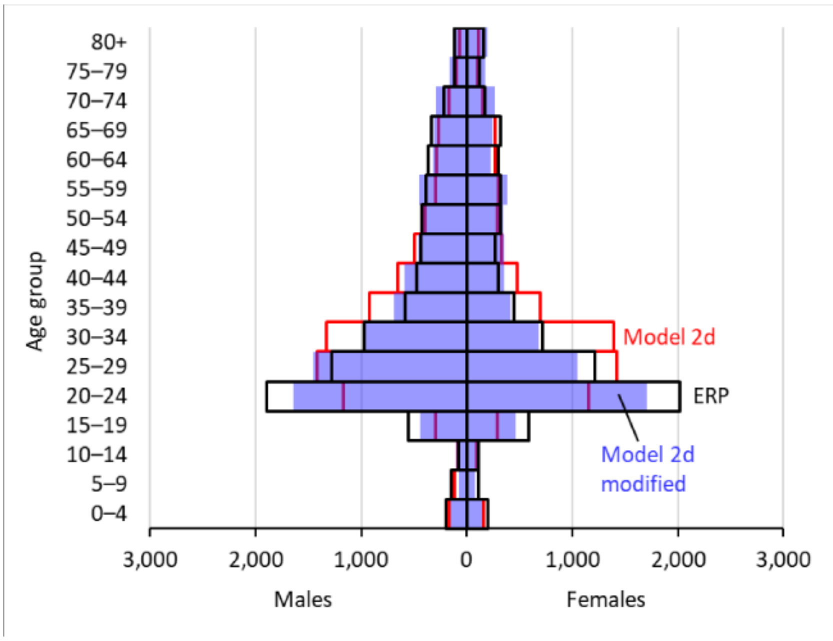

Figure 4: The population of Adelaide SA2 in 2016 projected from 2006

Notes: Model 2d is the constrained Hamilton-Perry CCR/CCD model with smoothing, Model 2d modified incorporates the No Change forecast for ages 0-4 to 60-64, and the ERP is the actual population.

Second, for a few areas, smoothing CCRs or CCDs over sex will yield poor forecasts. This is especially the case for small areas dominated by prisons or sex-specific boarding schools, or where there is male-dominated employment (e.g., in mining and construction). For these areas, the smoothing of CCRs and CCDs should not be sex-specific.

A third potential modification is to use alternative SA2 total projected population constraints. If reliable data are available on likely dwelling growth in urban areas, a housing-unit model could be employed to produce alternative population totals for those areas earmarked for substantial residential development. Similarly, alternative population totals might be based on other major developments, such as the opening of a new prison, or the closure of a remote mine accompanied by a big drop in the local town's population.

Fourth, where alternative total population forecasts of strong growth are chosen, different age profiles of CCRs and CCDs could be considered. Areas which grew little in the base period, but are projected to experience rapid dwelling growth, are likely to grow strongly in the younger adult ages at which migration is generally the highest. For these areas, the base period CCRs and CCDs could be substituted for those from similar areas which grew rapidly in the base period.

Fifth, it is best to avoid undertaking forecasts for the very smallest populations if at all possible. The largest median errors in forecasting the age-sex distribution of the population occurred for SA2 areas in the smallest population size category (Table 6). Where possible, it would be sensible to aggregate the smallest SA2 areas with one or two neighbouring areas to increase the chances of a more robust age-sex forecast. Although there is no obvious minimum population size cut-off, we recommend a population of at least 5,000 for obtaining useful forecasts.

A sixth suggestion is to create a 'test run' set of forecasts for the recent past, as we have done here. A comparison with actual populations will reveal if there are any problems with input data, and which areas require modifications. Finally, it is useful to undertake a careful evaluation of forecasts for plausibility and consistency prior to publication and distribution. A checklist of suggested items to assess is available elsewhere (author reference).

4.4. Concluding remarks

The main contribution of this study is to show that the alternative Hamilton-Perry CCR/CCD model using smoothed CCR and CCD age profiles and constrained to independent forecasts can provide a reasonable set of short-term age-sex forecasts for the majority of areas. We believe it is best used in situations when a full cohort-component model cannot be implemented, and with the modifications suggested above. The model out-performs the standard Hamilton-Perry CCR version, and is more accurate than the No Change model constrained to independent forecasts, especially at the older ages. Importantly, it requires little input data and input data preparation, allowing small area age-sex forecasts to be produced relatively quickly and easily, and for little cost.

Statements and declarations

Availability of data

The input data needed to replicate the population forecasts reported in this paper are available at [removed for peer review]

Competing Interests

We have no competing interests to declare.

Funding

[Removed for peer review]

References

ABS (2008). Population Projections, Australia, 2006 to 2101. https://www.abs.gov.au/AUSSTATS/abs@.nsf/allprimarymainfeatures/FF590BD57662A60BCA257C2E00172359?opendocument. Accessed 24 September 2021.

ABS (2018). Census of Population and Housing: Details of Overcount and Undercount, Australia, 2016. https://www.abs.gov.au/AUSSTATS/abs@.nsf/Lookup/2940.0Main+Features12016?OpenDocument. Accessed 5 October 2021.

Baker, J., Ruan, X., Alcantara, A., Jones, T., Watkins, K., McDaniel, M., Frey, M., Crouse, N., Rajbhandari, R., Morehouse, J., Sanchez, J., Inglis, M., Baros, S., Penman, S., Morrison, S., Budge, T., & Stallcup, W. (2008). Density-dependence in urban housing unit growth: An evaluation of the Pearl-Reed model for predicting housing unit stock at the census tract level. *Journal of Economic and Social Measurement, 33*(2-3), 155-163. https://doi.org/10.3233/JEM-2008-0301

Baker, J., Swanson, D. A., Tayman, J., & Tedrow, L. M. (2017). *Cohort Change Ratios and their Applications*. Springer International Publishing.

Baker, J., Swanson, D., & Tayman, J. (2020). The Accuracy of Hamilton–Perry Population Projections for Census Tracts in the United States. *Population Research and Policy Review*, 1-14. https://doi.org/10.1007/s11113-020-09601-y

de Beer, J. (2012). Smoothing and projecting age-specific probabilities of death by TOPALS. *Demographic Research, 27*(20), 543-592. https://doi.org/10.4054/DemRes.2012.27.20

Hamilton, C. H., & Perry, J. (1962). A short method for projecting population by age from one decennial census to another. *Social Forces, 41*(2), 163-170. https://doi.org/10.2307/2573607

Hauer, M. E. (2019). Population projections for US counties by age, sex, and race controlled to shared socioeconomic pathway. *Scientific Data, 6*(1), 1-15. https://doi.org/10.1038/sdata.2019.5

Inoue, T. (2017). A new method for estimating small area demographics and its application to long-term population projection. In D. A. Swanson (Ed.), *The Frontiers of Applied Demography* (pp. 473-489). Springer. https://doi.org/10.1007/978-3-319-43329-5_22

Rayer, S. (2008). Population Forecast Errors: A Primer for Planners. *Journal of Planning Education and Research*, 27(4), 417-430. https://doi.org/10.1177/0739456x07313925

Reinhold, M., & Thomsen, S. L. (2015). Subnational population projections by age: An evaluation of combined forecast techniques. *Population Research and Policy Review, 34*(4), 593-613. https://doi.org/10.1007/s11113-015-9362-0

Rowland, D. T. (2003). *Demographic Methods and Concepts*. Oxford: Oxford University Press.

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

Smith, S. K., Tayman, J., & Swanson, D. A. (2013). *A Practitioner's Guide to State and Local Population Projections*. Springer.

Swanson, D. A., Schlottmann, A., & Schmidt, B. (2010). Forecasting the population of census tracts by age and sex: An example of the Hamilton–Perry method in action. *Population Research and Policy Review, 29*(1), 47-63. https://doi.org/10.1007/s11113-009-9144-7

Tayman, J. (2011). Assessing uncertainty in small area forecasts: State of the practice and implementation strategy. *Population Research and Policy Review, 30*(5), 781-800. https://doi.org/10.1007/s11113-011-9210-9

Tayman, J., Smith, S.K. & Rayer, S. (2011). Evaluating population forecast accuracy: a regression approach using county data. *Population Research and Policy Review 30*(2), 235-262. https://doi.org/10.1007/s11113-010-9187-9

Tayman, J., & Swanson, D. A. (2017). Using modified cohort change and child-woman ratios in the Hamilton–Perry forecasting method. *Journal of Population Research, 34*(3), 209-231. https://doi.org/10.1007/s12546-017-9190-7

Tayman, J., Swanson, D. A., & Baker, J. (2021). Using synthetic adjustments and controlling to improve county population forecasts from the Hamilton-Perry method. *Population Research and Policy Review*, https://doi.org/10.1007/s11113-021-09646-7.

Wilson, T. (2015). New evaluations of simple models for small area population forecasts. *Population, Space and Place, 21*(4), 335-353. https://doi.org/10.1002/psp.1847

Wilson, T. (2016). Evaluation of alternative cohort-component models for local area population forecasts. *Population Research Policy Review, 35*(2), 241-261. https://doi.org/10.1007/s11113-015-9380-y

Wilson, T. (2017). A checklist for reviewing draft population projections. Charles Darwin University. https://researchers.cdu.edu.au/en/publications/a-checklist-for-reviewing-draft-population-projections

Wilson, T., Brokensha, H., Rowe, F., & Simpson, L. (2018). Insights from the evaluation of past local area population forecasts. *Population Research Policy Review, 37*(1), 137-155. https://doi.org/10.1007/s11113-017-9450-4

Wilson, T., Grossman, I., Alexander, M., Rees, P., & Temple, J. (2021). Methods for small area population forecasts: state-of-the-art and research needs. *Population Research and Policy Review*, https://doi.org/10.1007/s11113-021-09671-6

Appendix

Table A1: Mean error and bias of SA2 total population forecasts

	Model type				
	Hamilton-Perry CCR		Hamilton-Perry CCR/CCD		No Change
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
MAPE (%)					
5 year horizon					
1. Unconstrained	7.1	7.1	6.0	6.1	7.9
2. Constrained to national forecasts & SA2 total population forecasts	5.3	5.3	5.3	5.3	5.3
10 year horizon					
1. Unconstrained	15.1	15.2	10.8	10.9	13.2
2. Constrained to national forecasts & SA2 total population forecasts	9.5	9.5	9.5	9.5	9.5
Mean PE (%)					
5 year horizon					
1. Unconstrained	-1.1	-1.0	-2.8	-2.8	-7.1
2. Constrained to national forecasts & SA2 total population forecasts	-0.6	-0.6	-0.6	-0.6	-0.6
10 year horizon					
1. Unconstrained	1.1	1.1	-4.4	-4.3	-11.6
2. Constrained to national forecasts & SA2 total population forecasts	0.4	0.4	0.4	0.4	0.4

Table A2: Mean error and bias of SA2 age-sex population forecasts

	Model type				
	Hamilton-Perry CCR		**Hamilton-Perry CCR/CCD**		**No Change**
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
MAPE (%)					
5 year horizon					
1. Unconstrained	13.9	12.9	12.5	11.7	15.9
2. Constrained to national forecasts & SA2 total population forecasts	13.0	12.0	12.5	11.6	14.1
10 year horizon					
1. Unconstrained	22.9	21.7	17.8	17.1	22.6
2. Constrained to national forecasts & SA2 total population forecasts	18.8	17.6	17.6	16.7	19.3
Mean PE (%)					
5 year horizon					
1. Unconstrained	0.0	0.1	-1.7	-1.5	-6.7
2. Constrained to national forecasts & SA2 total population forecasts	0.4	0.4	0.5	0.6	0.6
10 year horizon					
1. Unconstrained	2.4	2.7	-3.0	-2.5	-11.3
2. Constrained to national forecasts & SA2 total population forecasts	1.8	2.0	1.9	2.2	2.1

Table A3: Errors in forecasting the population age-sex distribution of SA2 areas as measured by the mean Index of Dissimilarity

	Model type				
	Hamilton-Perry CCR		**Hamilton-Perry CCR/CCD**		**No Change**
Constraining	**Unsmoothed**	**Smoothed**	**Unsmoothed**	**Smoothed**	
5 year horizon					
1. Unconstrained	0.053	0.048	0.050	0.046	0.064
2. Constrained to national forecasts & SA2 total population forecasts	0.052	0.047	0.049	0.045	0.057
10 year horizon					
1. Unconstrained	0.074	0.067	0.067	0.062	0.087
2. Constrained to national forecasts & SA2 total population forecasts	0.072	0.066	0.065	0.061	0.075

Table A4: Errors in forecasting the population age-sex distribution of SA2 areas by population size and growth rate categories as measured by the mean Index of Dissimilarity, Model 2d

Jump-off year population size							
	<3,000	3,000-4,999	5,000-6,749	7,500-9,999	10,000-14,999	15,000-19,999	20,000+
5 year horizon	0.0834	0.0601	0.0475	0.0415	0.0368	0.0316	0.0290
10 year horizon	0.1039	0.0774	0.0630	0.0559	0.0506	0.0472	0.0446
Base period annual average growth rate (%)							
	<-1	-1 to 0	0 to 1	1-2	2-3	3-4	4+
5 year horizon	0.0580	0.0436	0.0418	0.0449	0.0457	0.0475	0.0503
10 year horizon	0.0742	0.0589	0.0559	0.0609	0.0621	0.0657	0.0684

Table A5: Regression results for the MedAPE over age and sex, Model 2d, 2016

Variable	Coefficient	Standard error	p value	% contribution†
Constant	6.577	3.563	0.0651	
2001-06 age-sex IoD	159.601	10.521	< 2.00e-16	36.2
Ln(jump-off population)	-1.106	0.335	0.000977	12.1
% different address	0.306	0.0347	< 2.00e-16	16.2
% Indigenous pop	0.0835	0.0176	2.30e-06	2.9
% communal dwellings	0.0413	0.0437	0.345	2.5
Base period growth rate	0.913	0.0633	< 2.00e-16	30.1
No. of observations	2,054			
Adjusted R-squared	0.335			

Notes. † Shapley decomposition of the contribution to R-squared