



DP2011/07

**Forecasting house price inflation: a model
combination approach**

Sarah Drought and Chris McDonald

October 2011

JEL classification: E17,E37

www.rbnz.govt.nz/research/discusspapers/

Discussion Paper Series

ISSN 1177-7567

DP2011/07

Forecasting house price inflation: a model combination approach*

Sarah Drought and Chris McDonald[†]

Abstract

In this paper we use a range of statistical models to forecast New Zealand house price inflation. We address the issue of model uncertainty by combining forecasts using weights based on out-of-sample forecast performance. We consider how the combined forecast for house prices performs relative to both the individual model forecasts and the Reserve Bank of New Zealand's house price forecasts. We find that the combination forecast is on par with the best of the models for most forecast horizons, and has produced lower root mean squared forecast errors than the Reserve Bank's forecasts.

* The Reserve Bank of New Zealand's discussion paper series is externally refereed. The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Reserve Bank of New Zealand. We thank Özer Karagedikli and participants at the 2010 New Zealand Association of Economists conference, 2011 New Zealand Econometric Society Group and various RBNZ seminars for helpful comments and discussion. We also acknowledge and thank Bernard Hodgetts and Aidan Yao for earlier work on some models.

[†] Address: Economics Department, Reserve Bank of New Zealand, 2 The Terrace, PO Box 2498, Wellington, New Zealand. *Email addresses:* Sarah.Drought@rbnz.govt.nz and Chris.Mcdonald@rbnz.govt.nz

1 Introduction

Forecasting house prices is important for monetary policy. This is because house prices can play a significant role in the evolution of business cycles. As such, house prices can be an important leading indicator of inflationary pressures.

One channel by which house prices can affect the macroeconomy is through consumption, via wealth and housing collateral effects. These effects were widely investigated following strong increases in house prices for many countries from early-mid 2000. For example, see Muellbauer (2007), Dvornak and Kohler (2007), Iacoviello and Neri (2010) and, for the New Zealand case, Hull (2003), and De Veirman and Dunstan (2008). Results from the literature suggest that increases in housing wealth can have significant positive effects on consumption. This effect has become stronger over recent years as financial deregulation increased households access to credit.

Developments in the housing market can also have an inflationary effect via private investment. As house prices increase relative to the value of housing-related construction costs (i.e. Tobin's q for residential investment), new housing becomes relatively more profitable. Moreover, as the value of assets that can be used as collateral increase (value of houses and land), the ability of individuals and firms to borrow and finance their investment increases. As a result, one would expect a strong positive correlation between house prices and private investment.

New Zealanders hold a large proportion of their wealth in housing assets.¹ Consequently, the evolution of house prices is particularly important for the New Zealand economy. Given this, there have been many studies that investigate 'fundamental' house prices in New Zealand, including Herring (2006) and Fraser et al (2008). Additionally, O'Donovan and Rae (1997), and Briggs and Ng (2009) model the short-run dynamics of New Zealand house prices around their fundamental level using error correction models. However, these studies are all done ex-post with no focus on out-of-sample forecasting. To our knowledge there is little published research that forecasts New Zealand house prices in real time.

However, there have been attempts to forecast house prices outside New Zealand using both structural and non-structural models. Das et al (2009) and Gupta et al (2009) use a variety of time-series models, including factor models and large scale Bayesian vector autoregression models, to forecast

¹ Hull (2003) reported this figure to be 80 percent.

house prices in South Africa and the United States respectively. In addition to a range of time-series models, Gupta et al (2011) use the 10 variable dynamic stochastic general equilibrium model of Iacoviello and Neri (2010) to forecast house prices in the United States.

In a similar manner, we use a range of time series models to produce quasi-realtime forecasts of house price inflation. Our models include an autoregression model (our benchmark model), single equation indicator models, Bayesian vector autoregressions, error correction models and a factor model, among others. Though we consider many models that cover a wide range of empirical macroeconomic modelling approaches and use many indicators, it is not an exhaustive list.

Estimating many models presents a problem about whether model selection (forecasting with the best model in real time for each horizon) or model combination would be the best forecast strategy. An advantage of model combination is that it helps to mitigate model uncertainty and may average out misspecification bias of individual models. Furthermore, a combination forecast may be more robust to unknown instabilities (structural breaks) than forecasts from individual models.

Accordingly, we consider two types of forecast combination: a simple average and a weighted average (where weights are based on inverse mean squared forecast errors) and compare the performance of these methods to a model selection approach. We find that the two combination approaches and the model selection approach produce similar root mean squared forecast errors for most forecast horizons. However, model selection outperforms the combination for forecasts between six and nine quarters ahead.

Lastly, we consider how our combined forecast for house prices performs relative to both the individual models and the Reserve Bank of New Zealand's (RBNZ) house price forecasts.² We find that the combination forecast always beats the autoregressive benchmark model and is on par with the best of the models at most horizons. The combination forecast has also produced lower root mean squared forecast errors than the RBNZ's published forecasts. As such, the combination approach for forecasting house prices would be a useful addition to the forecast process.

The remainder of the paper is structured as follows. Section 2 discusses model

² We use the RBNZ forecasts as a comparison because to our knowledge there are no house price forecasts published by an alternative source each quarter. While the RBNZ do not publish their house price forecasts on a regular basis, forecasts are produced each quarter for internal use.

combination. Section 3 introduces the eleven models used in the model suite. Section 4 outlines the data used to produce quasi-realtime forecasts. Section 5 outlines model evaluation and combination. Section 6 presents results and, finally, section 7 concludes.

2 Model Combination

Decision makers often have multiple forecasts of a particular variable available to them. Ex-ante it can be difficult to know how to best exploit the information contained in the individual forecasts. The two obvious options are to select the ‘best’ individual forecast or to combine the individual forecasts to produce a single summary forecast. The latter is known as model (or forecast) combination and has been widely used in the forecast literature, with empirical applications dating back at least to Bates and Granger (1969).

Timmermann (2006) provides a recent survey of the large literature on forecast combination. Timmerman outlines four reasons for using forecast combinations rather than relying on the forecast from the ex-ante best individual model. Two of these relate to our motivation for using a forecast combination approach.

Firstly, a combination forecast may be more robust to the misspecification biases of individual models and measurement errors in the underlying datasets. Because the true data generating process is unknown and likely to be more complex than the most flexible and general model used by the forecaster, the individual models will be subject to an unknown misspecification (omitted variables) bias. A combination will be useful in the case where the individual models are subject to different bias. In this case it is plausible that the combination will average out the biases and improve forecast accuracy.

Secondly, a combination forecast may be more robust to unknown instabilities (structural breaks) than forecasts from an individual model. This is because individual models may be affected differently by structural breaks with some models adapting quickly while others are slow to adjust (parameters update slowly to post-break data). Structural breaks are generally difficult to detect in real time so a combination can be useful because the resulting forecast will contain information from various models with different degrees of adaptability. Thus, if a break was to occur the combination forecast would be influenced by some models whose coefficients adjust quickly to reflect the post-break data.

Empirical studies have shown that combination forecasts frequently outperform forecasts from the best-performing model in real time. Timmermann (2006) highlights that simple combination schemes (ones that do not require the estimation of many parameters) are hard to beat. For example, a combination using equal weights or weights based on past forecast performance, such as inverse mean squared error weights, tend to produce more accurate forecasts than a combination using optimal weighting schemes based on the full variance-covariance matrix of forecast errors (See Timmermann (2006) and the references therein).

For some of the reasons outlined above, forecast combination is becoming increasingly popular in central banks. The Reserve Bank of New Zealand currently uses a suite of empirical models to produce forecasts for key macroeconomic variables. The combined forecasts are used to highlight any risks around the central projection.³ Other central banks that use forecast combination include the Bank of England, the Riksbank (Sweden), Norges Bank (Norway) and the Bank of Canada, see Bjørnland et al (2009) for an overview.

3 Models

To forecast house prices, we use eleven time series models with a number of different indicators. We start with models as simple as an autoregressive process and single equations, and then discuss more dynamic models - some of which are fairly data rich. While we have tried to consider a broad range of models that cover a wide range of empirical macroeconomic modelling approaches, it is not an exhaustive list.

All models are estimated in quasi-realtime producing forecasts up to 10 quarters ahead for the Quotable Value New Zealand (QVNZ) nominal house price index from 1994Q1 to 2010Q4. That is, each model is estimated at every forecast date using real time data or quasi-realtime data.⁴ Most models are estimated on data dating from 1988Q1 with any exceptions to this noted. This start date is influenced by data limitations and also the economic reforms in the 1980s. While there were still changes after this date (for example, the adoption of the inflation target in 1990), we did not want to shorten the

³ Combination forecasts are produced for GDP, consumption, CPI, tradable CPI, non-tradable CPI, the 90-day interest rate and the NZD TWI.

⁴ We define quasi-realtime data as the most recent vintage of the data cropped back to what would have been available at the time of the forecast.

length of the time series any further. Section 4 contains more information about the data.

Autoregression (AR)

We started with a simple autoregressive model to forecast quarterly QVNZ house price inflation ($HPinf_t$). This type of model has been shown to forecast reasonably well and we use this as our benchmark model. The number of lags (p) ranges from one to four and is chosen using the Bayesian information criteria. The following equation shows an AR(p) model:

$$HPinf_t = \gamma_0 + \gamma_1 HPinf_{t-1} + \dots + \gamma_p HPinf_{t-p} + \epsilon_t \quad (1)$$

Migration and Mortgage Rate indicator (MM indicator)

Migration flows and mortgage rates are important factors of the demand for housing. Incoming migrants need houses and out-going migrants vacate them. This affects housing activity through both house prices and construction activity. Coleman and Landon-Lane (2007) show that the relationship between migration flows and house prices was particularly strong in New Zealand over the past half century. Mortgage rates determine the cost of financing a house purchase and consequently are also an important determinant of housing demand.

To capture these relationships, we include permanent and long term (PLT) arrivals, PLT departures, and the 5-year mortgage rate in a single equation model. The equation is estimated using OLS with the independent variables all lagged by two quarters. This lag was chosen because it corresponds to the tightest in-sample correlation with house prices.

$$HPinf_t = \gamma_0 + \gamma_1 PLTA_{t-2} + \gamma_2 PLTD_{t-2} + \gamma_3 R_{t-2} + \epsilon_t \quad (2)$$

To forecast house price inflation with this equation we need forecasts of PLT arrivals, PLT departures, and the 5-year mortgage rate. These forecasts are generated using an AR process.

Bayesian Vector Autoregression (Small BVAR)

This model includes the same four variables as the MM indicator: quarterly house price inflation, PLT arrivals, PLT departures, and the 5-year mortgage rate, and models them in VAR. The VAR framework is beneficial because all variables are treated endogenously so forecasts for each variable are done within the model.

Consider a VAR(p) model:

$$Y_t = c + B_1 Y_{t-1} + B_2 Y_{t-2} + \dots B_p Y_{t-p} + v_t \quad (3)$$

where the Bayesian information criteria is used to select the optimal number of lags (p), with p ranging between one and four.

We estimate the VAR using Bayesian techniques, in part because of few degrees of freedom, though later in the sample the number of data observations was ample. We use Minnesota priors that assume the mean of prior distribution is a random walk. The specification of the standard deviation of the prior imposed on variable j in equation i at lag k is:

$$\sigma_{ijk} = \theta w(i, j) k^{-\phi} \left(\frac{\hat{\sigma}_{uj}}{\hat{\sigma}_{ui}} \right) \quad (4)$$

Where θ is the ‘overall tightness’ parameter, reflecting the standard deviation of the prior on the first lag of the dependent variable. The $k^{-\phi}$ term is the lag decay parameter. Increasing ϕ reduces the standard deviation of the priors on lags greater than one, imposing the belief that more recent lags contain more useful information than more distant ones. Also, $w(i, j)$ allows us to weight the priors on variables differently. For a good summary of the priors see LeSage (1999).⁵⁶

⁵ We use fairly typical hyperparameter values, θ is set to 0.2 and ϕ is one. We place more weight on own lags of PLT arrivals and PLT departures in their corresponding equations than on lags of other variables using $w(i, j)$. This reflects our belief that the movements in migration flows are unlikely to be well explained by mortgage rates or house prices.

⁶ We use the LeSage MATLAB package to estimate this Bayesian VAR, using Gibbs sampling. For further details on these functions see LeSage (1999).

Bayesian Vector Autoregression (BVAR)

Using a similar Bayesian VAR technique to that above, we take the four variables discussed for the small BVAR and MM indicator models and add seven additional variables. We add these variables to increase the number of possible indicators for house prices. Notably, we aim to include variables that have some leading information, or are key macroeconomic indicators. The additional variables are: quarterly consumer price inflation, the quarterly difference in the unemployment rate, quarterly residential investment growth, quarterly GDP growth, the terms of trade, the New Zealand dollar TWI, and the quarterly difference in the Australian unemployment rate.

In this 11 variable VAR, the number of parameters gets very large. As our dataset only starts in 1988, we have few degrees of freedom available to estimate the parameters without over-fitting. We use Bayesian shrinkage to overcome this problem. To set priors we adopt the method used by Bloor and Matheson (2008) which is similar to the algorithm used by Bańbura et al (2008). Using this method we firstly estimate a baseline VAR (the small BVAR) using OLS (the equivalent of Bayesian estimation with very loose priors). Then, the overall tightness hyper-parameter, θ (reflecting the standard deviation of the prior on the first lag of the dependent variable), on the larger BVAR is tightened so that the in-sample fit is equal to the in-sample fit of the baseline VAR.⁷ The number of lags in the large BVAR is determined using the Bayesian information criteria (the same number of lags is used in the baseline BVAR).

REINZ Monthly Vector Autoregression (REINZ VAR)

The Real Estate Institute of New Zealand (REINZ) releases monthly housing data including a house price index, the number of house sales, and the median days to sell. We use these three series in a monthly VAR to capture the most recent movements in the housing market. We also include the 90-day interest rate to allow for changes in monetary policy.⁸

We estimate this VAR using Bayesian techniques, the same as those used to

⁷ As in the small BVAR, the decay hyper-parameter (ϕ) is set to one. We also use the weighting matrix (w) to put more weight on own lags of the Australian unemployment rate and on the terms of trade in their corresponding equations than on the other variables. We argue that the New Zealand economy is too small to have a significant impact on international prices or the Australian unemployment rate.

⁸ We use data starting in 1992M1, as this is when the REINZ data set begins.

estimate the small BVAR.⁹ Again, we use the Bayesian information criteria to choose the lag length, with the lags ranging between one and six.¹⁰ This VAR forecasts the REINZ variables at a monthly frequency. These forecasts are collapsed to a quarterly frequency by taking the quarterly average.

In a second step, we forecast quarterly QVNZ house price inflation (the house price series forecast in other models) using a bridging equation. We use OLS to regress quarterly QVNZ house price inflation ($HPinf$) on quarterly REINZ housing price inflation ($RHPinf$) and house sales (HS). We include house sales in the equation because this improves the in-sample fit between the QVNZ house price inflation and REINZ house price inflation.

$$HPinf_t = \gamma_0 + \gamma_1 RHPinf_t + \gamma_2 HS_t + \epsilon_t \quad (5)$$

Using the forecasts for the REINZ variables from the first stage, equation 5 can be used to generate forecasts for quarterly QVNZ house price inflation.

Error Correction Model (ECM)

We use a simple two-step error correction model to forecast nominal house prices. The first step estimates the fundamental house price level. The second step uses the error correction term (deviation of house prices from their fundamental level), amongst other data, to forecast the growth in house prices. This model is related to work by O'Donovan and Rae (1997), the International Monetary Fund (2003), Abelson et al (2005) and Briggs and Ng (2009), among others.

In the long-run equation, the fundamental house price level is determined by nominal GDP, the working age population and the user cost of capital.¹¹ Nominal GDP is designed to capture the income effect on housing demand. We include the working age population to capture any influence that a changing population will have on housing demand.¹² For the user cost of capital we use the floating first-mortgage new customer interest rate less the

⁹ The tightness hyper-parameter (θ) is set to 0.1. The decay hyper-parameter (ϕ) is 1. We treat all variables the same in every equation by making $w(i, j) = 1$ for all i and j .

¹⁰ Prior to 1996, we restrict the VAR to just one lag due to a very short sample.

¹¹ This specification is similar to that used in International Monetary Fund (2003) and Briggs and Ng (2009), however our model also includes a population variable and the user cost of capital rather than the mortgage rate.

¹² We use the working age population rather than total population to better reflect the cohort of the population who will be house buyers.

expected capital gain from housing, which we proxy by the most recent 3 year moving average of nominal annual house price inflation.

Long-run equation:

$$hp_t = \alpha_0 + \alpha_1 gdp_t + \alpha_2 pop_t + \alpha_3 UC_t + \epsilon_t \quad (6)$$

where hp is the log of the QVNZ house price index, gdp is the log of nominal GDP, pop is the log of the working age population, and UC is the user cost.¹³

The long-run equation is estimated using OLS on data dating from 1970Q1. Initial testing of the data with the end date ranging from 1994Q1 to 2010Q4 suggests that the log levels of nominal house prices, nominal GDP and the working age population all contain unit roots. We also find at least one co-integrating vector among the variables included in the long-run equation. This suggests that an error correction model is appropriate.

The short-run equation uses the lag of the error correction term (ϵ_{t-1}) to forecast house price inflation. The lag of house price inflation and the working age population growth are also included in this equation.¹⁴

Short-run equation:

$$\Delta hp_t = \beta_0 + \beta_1 \epsilon_{t-1} + \beta_2 \Delta hp_{t-1} + \beta_3 \Delta pop_{t-1} + \varsigma_t \quad (7)$$

The short-run equation is also estimated using OLS.¹⁵ The coefficients on variables are generally correctly signed for the different vintages and are highly significant (at the 1 percent level in many cases).¹⁶ Furthermore, we can reject the hypothesis that the error correction term contains a unit root. Thus, the error correction mechanism will help house prices adjust to the model-determined ‘fundamental’ level over time.

¹³ The user cost has not been logged because there may have been periods over our sample period where this was negative.

¹⁴ The short-run equation is estimated over a shorter sample dating from 1988Q1 to be more consistent with the sample used for the other models.

¹⁵ To generate forecasts for house price inflation this model requires forecasts for the independent variables. Nominal GDP, working age population, and the mortgage rate are all forecast using an AR(1) process.

¹⁶ For example, in the long run equation α_1 and α_2 are positive, whereas α_3 is negative. In the short run equation, β_1 is negative, while β_2 and β_3 are positive.

Factor model

The most data-rich model in our suite is a dynamic factor model. This model uses factors from 21 data series, including general macroeconomic data, interest rate data and housing specific data, such as consents, all dating from 1988.¹⁷ This is typically a small number of series to be included in a factor model and a large Bayesian VAR, similar to the BVAR discussed earlier, could have been used instead. However, to add another dimension to our model suite we opted to use a factor model. Furthermore, Boivin and Ng (2006) find that extracting factors from an increasingly large number of series is not necessarily optimal. Thus, we include a carefully selected set of indicators and avoid additional series that may introduce more noise into the factor estimation while providing little additional information about the housing market.

The forecasts are made using the following equation, where $HPinf_{t+h}$ is the quarterly house price inflation forecast at horizon h , f_t is a matrix of factors and (L) denotes the variations on lag length. Forecasts from this model are estimated using direct forecasting methods rather than iterative methods used in the other models.

$$HPinf_{t+h} = \phi + HPinf_t + \gamma(L)f_t + \vartheta_{t+h} \quad (8)$$

This factor model allows for many variations in its structure, similar to Stock and Watson (2002). In particular, it allows for variations in the number of factors and lags of the factors.¹⁸ The Bayesian information criteria is used to choose between different structures.

Rental yield model

Here we apply a simple asset price model to the New Zealand housing market. We use the framework outlined by Weekin (2004), who applies this model to the UK housing market. This asset price model is one where households either rent or own houses, which are seen as durable assets providing a flow of housing services. According to this model, in equilibrium, real house prices Ph will be given by:

¹⁷ The series included are listed in appendix A.

¹⁸ To ensure parsimony, the maximum number of factors and lags of factors is set to two.

$$Ph = \frac{D}{r_f + \rho - g} \quad (9)$$

where D is the current real rental less maintenance and other non-interest costs, r_f is the risk free real interest rate, ρ is the housing risk premium and g is the expected long-run growth rate of rentals.¹⁹

The intuition behind this framework is quite simple. The equilibrium price is the discounted value of the expected future net real rental flow.²⁰ Changes in either the net real rental flow or the discount rate will motivate portfolio shifts into or out of housing, thus affecting the equilibrium price.

Equation 9 calculates an estimate of the fundamental level of real house prices. The deviation of house prices from their fundamental level is then used in a VAR that also includes quarterly real house price inflation, PLT arrivals, and PLT departures.²¹ We estimate this VAR using Bayesian estimation with standard Minnesota priors, the same as those used in the small BVAR. The number of lags included in the VAR can range from one to four and are chosen using the Bayesian information criteria.

This model is estimated with data starting in 1985Q1.²² The VAR is used to forecast all the variables except the deviation of house prices from their fundamental level. The forecast for the deviation is found by first forecasting the ‘fundamental’ price level for the next horizon using equation 9 and then taking the difference between this forecast and the 1-horizon ahead forecast for house prices from the VAR.²³ The forecast deviation of house prices from

¹⁹ Following DTZ New Zealand (2004), we assume net rentals to be 75 percent of the gross rent and deflate the annual net rental series by the CPI to arrive at an estimate of the net real rental. The expected growth in real rentals is not observable so we assume it equal to the long-run average growth rate in real rentals over the past two decades (about 1 per cent per annum).

²⁰ The discount rate is the risk free real interest rate plus the risk premium households require for holding houses rather than risk free assets (government bonds). Because this risk premium is unobservable (and likely to vary through time) we use the lowest mortgage rate available (from floating, 2-year and 5-year rates) less inflation as a proxy for the discount rate. To smooth out the quarter to quarter noise we use the HP-trend of the discount rate series with a smoothing parameter of 5000.

²¹ Real house prices are calculated by deflating nominal house prices with the CPI.

²² A longer sample is used for this model to better capture the long-run relationships that may exist for the ‘fundamental’ house price level in this context. However, we could not extend the sample further due to data limitations.

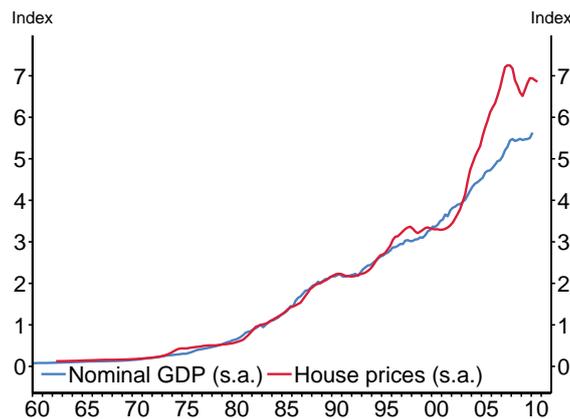
²³ This process requires forecasts for the minimum real mortgage rate, CPI inflation and real rents. These variables are forecast using an AR(4) model estimated on the most recent 10 years of data.

their fundamental level is then used in the VAR to forecast house prices for the next horizon, and the process repeats for all forecast horizons. Finally, the forecast of CPI inflation is added to the real house price inflation forecast to produce a forecast for nominal house price inflation.

Nominal GDP error correction model (YECM)

Over the long term, New Zealand’s nominal GDP and house prices have tended to move together (figure 1). Briggs and Ng (2009) discuss how this might reflect a fixed supply of land. They suggest that as real aggregate income rises over time, the real price of land would rise at a similar rate.

Figure 1
House price index and nominal GDP



We use a error correction model to forecast house prices while accounting for this long run relationship. We do this in a two stage process, first we estimate a bivariate regression to find the long run relationship between the log of house prices (hp_t) and the log of nominal GDP (gdp_t), shown in equation 10. This equation is estimated using OLS on data dating back to 1970, a long period to best capture the long-run relationship.

$$hp_t = \gamma_0 + \gamma_1 gdp_t + \epsilon_t \quad (10)$$

In the second stage of this model, the error term from equation 10 (ϵ_t) is added to a VAR that includes quarterly house price inflation, the 5-year mortgage rate, PLT arrivals and PLT departures, similar to the process of

the rental yield model. Again, we use Bayesian methods to estimate the VAR.²⁴ Standard priors are used (the same as used in the small BVAR), and the Bayesian information criteria is used to determine the optimal number of lags between one and four.

As in the rental yield model, all variables are forecast in the VAR except for the error correction term (ϵ_{t+h}). The forecast for the error correction term is the difference between the house price forecast (from the VAR) and the forecast of the model-implied fundamental level (found using equation 10, where GDP is forecast forward using an AR process with four lags). The forecast for ϵ_{t+h} is then used in the VAR to generate forecast for the next period, and the process repeats for all forecast horizons.

Regional BVAR

Regional house prices do not necessarily move together. Figure 2 shows the growth in house prices for New Zealand’s three main centres, Auckland, Wellington and Christchurch.²⁵ House prices throughout New Zealand may be driven differently by changes in migration and mortgage rates. For example, if a large proportion of migrants entering New Zealand migrate to Auckland, then Auckland’s house prices might react more strongly to a spike in PLT arrivals. Alternatively, if more properties in Wellington were owner-occupied then a rise in mortgage rates might have a smaller impact on house prices, relative to a market where investors were more prominent.

To model and forecast house prices in each centre, we use the same framework as used in the small BVAR. Each centre faces the same mortgage rate and immigration series (we do not have access to regional immigration data). Therefore, the regional VARs differ only by the house price series included and the parameter estimates. For each VAR, we use Bayesian estimation with Gibbs sampling.²⁶ We allow the Bayesian information criteria to select the number of lags, ranging from one to four.

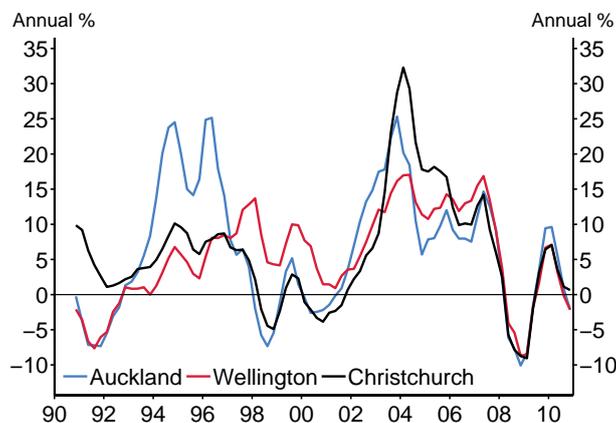
The forecasts from each centre are aggregated using a contemporaneous bridging equation that links the three regional house price series to the

²⁴ The VAR is estimated on data starting in 1988Q1.

²⁵ The regional house price series are also published by Quotable Value New Zealand and are available from 1990Q1.

²⁶ The priors are the same used for the small BVAR: the overall tightness hyperparameter, θ is set to 0.2, decay parameter, ϕ is set to one, and the weighting matrix (w) treats the migration variables as exogenous.

Figure 2
Regional house price inflation



national house price series. We allow for persistent differences between the weighted average and the official national house price series by incorporating a lag of the official series in the equation.

A model of supply and demand for housing

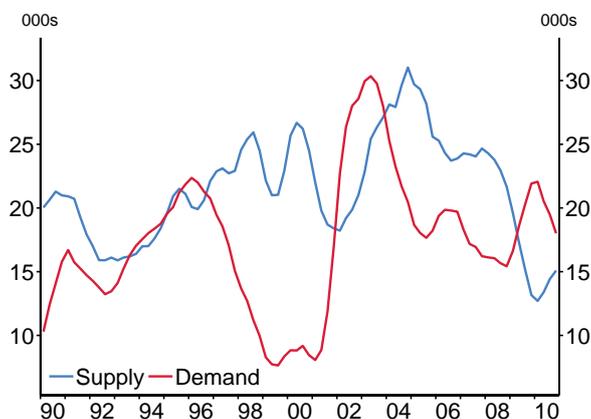
This model forecasts house prices given changes in supply and demand for housing. To do this, two data series are created: the number of houses needed (demand) and the number of houses built (supply).

The demand for housing is approximated by the quarterly increase in the population divided by the number of people per household. The number of people per household is estimated to be equal to total population divided by the stock of residential dwellings. The supply of housing is approximated by the change in the stock of residential dwellings (figure 3).²⁷

The supply and demand series are used in a VAR that also includes the 5-year mortgage rate and quarterly house price inflation. The VAR is estimated using Bayesian techniques with Gibbs sampling. Again, the number of lags can range between one and four and is determined using the Bayesian information criteria. The priors use the same hyperparameters as the small

²⁷ We recognise that supply tends to be larger than demand. This reflects that we have not accounted for houses that are destroyed, which we have little information about.

Figure 3
Annual supply and demand for houses



BVAR. However, there are no adjustments to the weighting matrix, w , so all variables are treated the same across different equations.

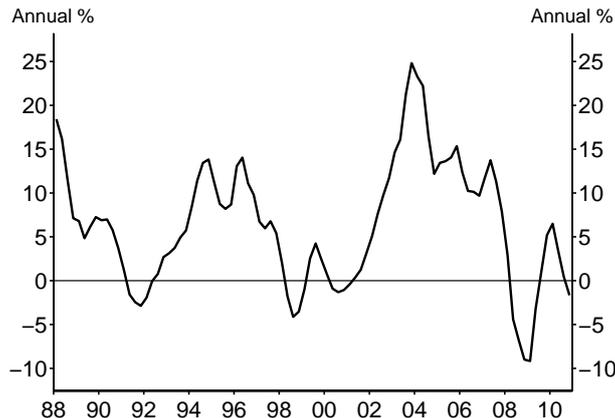
4 Data

Where possible we use real time data to produce the historic house price forecasts. That is, the data available at the time the forecast would have been made. For key variables used in the RBNZ forecast process (such as GDP, CPI and house prices), real time data has been stored when the RBNZ forecasts are published each quarter. However, for many other variables real time data is not available. For these variables we use the most recent data and crop it back to what would have been available at the time of the forecast (quasi-realtime). In New Zealand, the real variables from the National Accounts are subject to the largest revisions. Generally, other data is not revised or the revisions are relatively small. Because we have real time data for most of the real variables that we use, it is our opinion that the use of quasi-realtime data for some variables would have little effect on the results presented in this paper.

The house price series that we forecast in this paper is the Quotable Value New Zealand (QVNZ) quarterly house price index, shown in figure 4. This series has been forecast at the RBNZ for more than a decade. However, the QVNZ index is somewhat untimely (four-month lag). As such, in much of

our analysis we use the REINZ monthly house price index, which only has a two week lag, to forecast the first quarter.²⁸

Figure 4
QVNZ house price series



We considered forecasts both with and without the REINZ monthly data. As this additional information was found to improve forecast performance, we have included this data where possible in our models.²⁹ One caveat to this approach is that the REINZ index was only developed in 2009 (but backdated to 1992) so comparing the forecast performance of the models to the RBNZ in the first quarter may not be entirely fair. However, in trying to develop accurate forecasting models it is important to consider all available information. Additionally, the information incorporated in the REINZ index has essentially been available to the RBNZ for a number of years.

The data series used in each model (with the exception of the factor model) were mentioned in section 3. In total, 39 data series from a range of sources are used in estimating the model suite. The individual series, grouped by data source, are outlined in appendix A.

5 Evaluation

Initially, all models are estimated on data up to 1994Q1 and are used to forecast the quarterly percent change (qpc) of house prices up to 10 quarters

²⁸ See McDonald and Smith (2009) for information on the REINZ house price index.

²⁹ The REINZ VAR is the only model which generates it's own first-quarter forecast. All other models use the REINZ index for their first forecast.

ahead.³⁰ All models are then re-estimated using new data each quarter and subsequently produce a new set of forecasts. This process is repeated until 2010Q4 which is the final estimation period.

The models are all evaluated for each forecast horizon using actual data. The relative performance of the individual models varies across forecast horizons (some models are better for near-term forecasting and others are better for medium-term forecasting), so the combination weights are horizon-specific.

5.1 Forecast combinations

We consider two combination approaches and a model-selection approach to produce point forecasts at each horizon. The combination approaches use equal weights and mean squared error weights. When using equal weights, each model receives a $1/N$ weighting, where N is the number of models. Mean squared error (MSE) weights are calculated using the inverse of the model's mean squared error. Hence, the model with the lowest average forecast error receives the highest weight for the combination forecast.

MSE weights:

$$MSE_{i,h} = \frac{1}{T} \sum_{t=1}^T (y_{t+h} - \hat{y}_{t+h})^2 \quad (11)$$

$$MSEweight_{i,h} = \frac{\frac{1}{MSE_{i,h}}}{\sum_{i=1}^N \left(\frac{1}{MSE_{i,h}}\right)} \quad (12)$$

where T refers to the number of forecasts h -horizons ahead and N to the number of models. Forecasts from 1994Q1 to 1997Q4 are used to initialise the MSE weights, allowing quasi-realtime combination forecasts to be calculated recursively from 1998Q1.

In addition to the two combination methods, we also consider a model selection approach. This method forecasts using the model with the lowest mean squared error at the forecast date (the ex-ante best model), for each horizon. The fact that a model has forecast accurately in the past, does not mean that this will be true in the future so the selected model may

³⁰ Because of the four-month lag in QVNZ house price outturns, the first two forecasts are 'nowcasts'. Specifically, the first forecast quarter forecast is the forecast for the previous quarter and the second quarter forecast is for the current quarter.

change over time. Model selection forecasts are also produced recursively from 1998Q1.

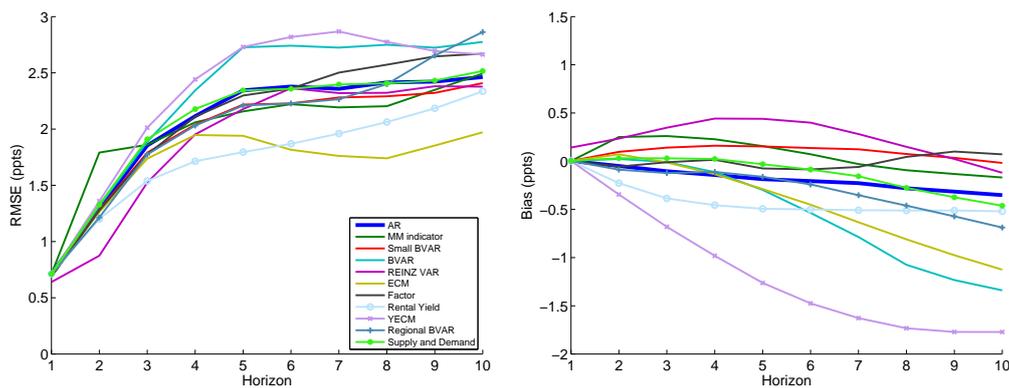
6 Results

In this section we describe the main results. First, we look at the forecast performance of the individual models. We then compare the performance of the two combination approaches outlined in section 5 against the model selection approach. Finally, our preferred combination forecast, which uses mean squared error (MSE) weights, is compared to both the individual model forecasts and the Reserve Bank of New Zealand’s house price forecasts. All results shown below are evaluated over the period 1998Q1 to 2010Q4.

6.1 Individual model results

Figure 5 shows the root mean squared errors (RMSEs) and biases of each model’s forecasts. Table 1 shows the RMSEs expressed relative to the benchmark AR model. The blue entries in this table are significantly better than the AR benchmark, while the red entries are significantly worse.³¹

Figure 5
Root mean squared errors (left) and biases (right) of individual forecasts



³¹ Statistical significance is measured using the Diebold and Mariano (1995) test at the 95 percent confidence level.

Table 1

Root mean squared errors of individual forecasts relative to AR benchmark

	1	2	3	4	5	6	7	8	9	10
AR (absolute RMSE)	0.71	1.30	1.85	2.11	2.34	2.38	2.36	2.41	2.42	2.46
MM indicator	1.00	1.38	1.00	0.97	0.92	0.93	0.93	0.91	0.97	1.01
Small BVAR	1.00	0.99	0.97	0.96	0.95	0.94	0.97	0.95	0.96	0.98
BVAR	1.00	1.02	1.02	1.11	1.16	1.15	1.16	1.14	1.12	1.13
REINZ VAR	0.90	0.67	0.82	0.92	0.93	0.99	0.98	0.96	0.98	0.97
ECM	1.00	0.97	0.94	0.92	0.83	0.76	0.75	0.72	0.77	0.80
Factor	1.00	0.99	0.96	1.00	0.98	0.99	1.06	1.07	1.09	1.08
Rental Yield	1.00	0.93	0.83	0.81	0.77	0.79	0.83	0.85	0.90	0.95
YECM	1.00	1.05	1.09	1.15	1.16	1.19	1.22	1.15	1.11	1.08
Regional BVAR	1.00	0.94	0.96	0.96	0.94	0.94	0.96	0.99	1.10	1.16
Supply and Demand	1.00	1.03	1.03	1.03	1.00	0.99	1.02	1.00	1.00	1.02

The AR model performs comparably to the other models, with the RMSE for all forecast horizons well within the range of RMSEs from the other models. However, there are a few models whose forecasts are considerably more accurate than the AR benchmark.

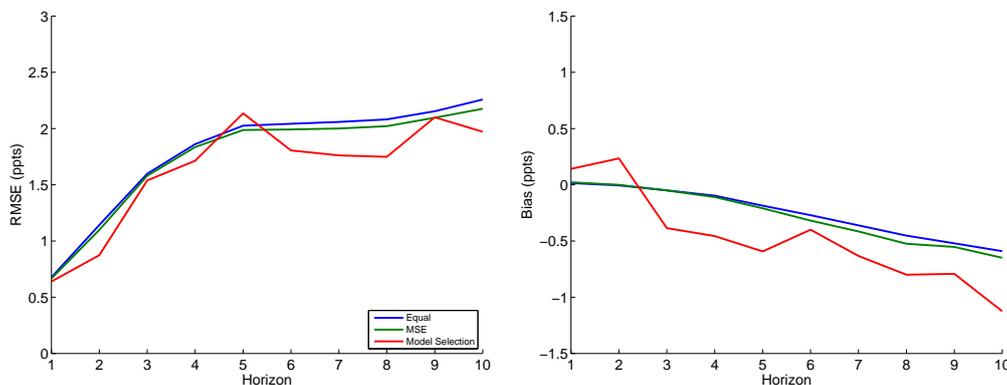
The REINZ VAR forecasts particularly well for the first three quarters, with the RMSE for the second quarter 33 percent lower than the RMSE from the AR model. The REINZ VAR uses monthly housing-specific data, whereas a majority of the data used by the other models is general macroeconomic data measured at a quarterly frequency. Thus, the use of higher frequency and more timely housing-specific data is beneficial for forecasting house price inflation in the near term.

For horizons beyond three quarters, the ECM, rental yield model and small BVAR model are the most accurate. The forecast performance of the rental yield model is statistically better than the AR model for forecasts three to eight quarters ahead, while the forecast performance of the small BVAR is statistically better than the AR model for forecasts six quarters ahead. These results suggest that factors such as income, population flows, rental yields and mortgage rates are important determinants of house prices over the medium term.

6.2 Combination results

We combine the individual forecasts using equal and MSE weights. Figure 6 shows the RMSEs and biases of these combination forecasts compared to the forecasts produced using model selection.

Figure 6
Root mean squared errors (left) and biases (right) of forecasts using different weighting methods



Over our sample period, forecasts from the best-performing model (model selection) have been more accurate than the combination forecasts for most horizons, particularly six to nine quarters ahead. The good performance of model selection over these horizons is likely due to the ECM’s forecasts being considerably more accurate than most of the other models. This result is at odds with the empirical evidence that combination forecasts frequently outperform forecasts from the best performing model in real time, outlined by Timmermann (2006).

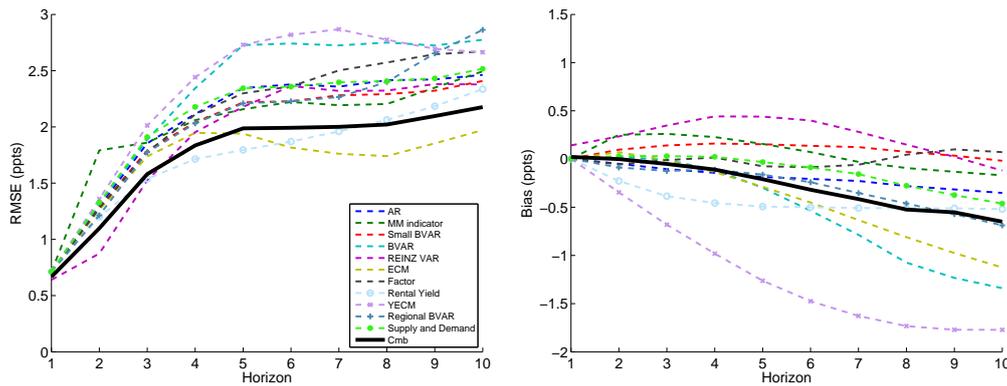
While this is a surprising result, the use of a combination forecast may still be beneficial for multiple reasons. First, the difference between the RMSEs of the combination and model selection forecasts is generally small, and the combination is less biased than the model selection forecast. Second, the good performance of a particular model in the past does not guarantee good forecast performance in the future. Finally, as outlined in section 2, a combination approach will likely be more resilient to structural breaks and misspecification biases than forecasts from an individual model.

The difference in forecast performance between the two combination methods is marginal. For analysis hereafter we focus on the MSE-weighted combination because this allows the weights on each model to vary over time. Consequently, in contrast to equal weights, MSE-weights should adapt somewhat to any structural breaks in the data.

The MSE-weighted combination has performed well relative to the individual models over our sample period. This is shown in figure 7 where the combina-

tion is shown by the solid black line. The RMSE of the combination is only slightly higher than the best individual model at each forecast horizon.

Figure 7
Root mean squared errors (left) and biases (right) of individual and combination forecasts



The combination forecast has a small bias compared to some of the individual models. One of the motivations for using a combination forecast was that it may average out biases from individual models. This appears to be the case here.

Do the combination forecasts outperform the Reserve Bank of New Zealand’s forecasts?

We now compare the performance of the MSE-weighted combination forecasts with the RBNZ’s forecasts.³² The RMSEs and biases of both forecasts are shown in figure 8, with the RMSEs also in table 2.³³ Over the evaluation period, the combination forecasts have been considerably more accurate than the RBNZ forecasts for all horizons. The improvement in forecast performance is statistically significant for forecasts up to nine quarters ahead. However, a comparison of these forecasts for the first quarter is not entirely fair because the REINZ series, used by most models for their first quarter forecast, was not developed until 2009.

³² The RBNZ forecasts are used as a comparison because to our knowledge there are no house price forecasts published by an alternative source each quarter.

³³ The blue entries in table 2 show the forecast horizons where the forecast performance of the combination has been significantly better than the RBNZ. Statistical significance is measured using the Diebold and Mariano (1995) test at the 95 percent confidence level.

Figure 8
Root mean squared errors (left) and biases (right) of combination and Reserve Bank forecasts

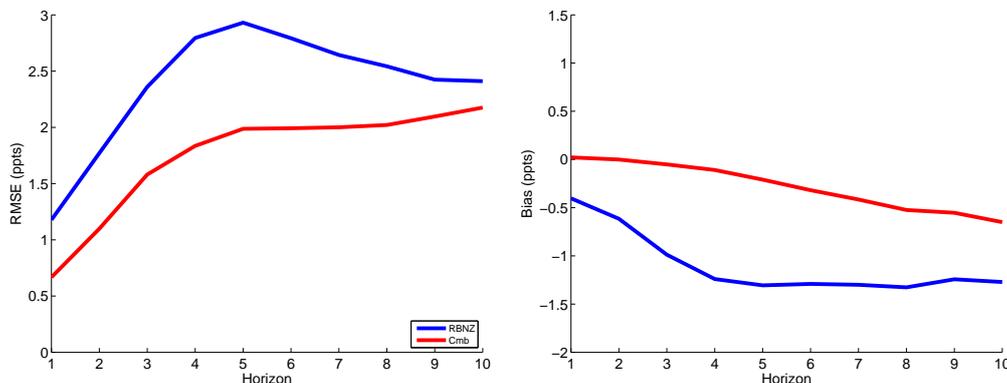


Table 2
Root mean squared errors

	1	2	3	4	5	6	7	8	9	10
RBNZ	1.18	1.77	2.36	2.79	2.93	2.79	2.64	2.54	2.42	2.41
Combination	0.67	1.10	1.58	1.84	1.99	1.99	2.00	2.02	2.10	2.18

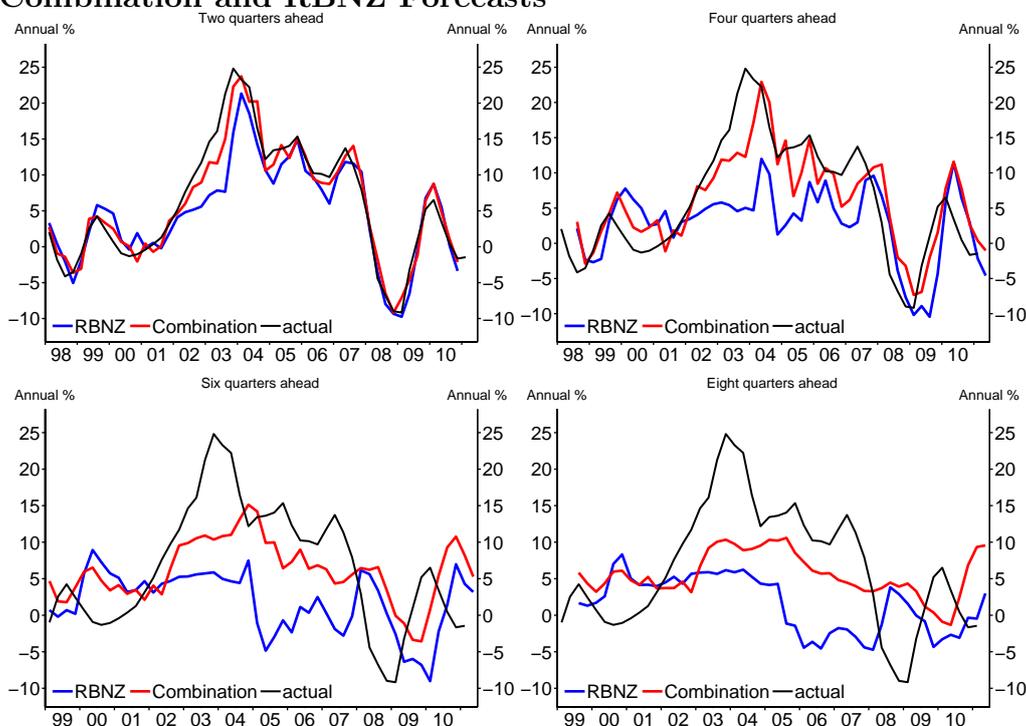
During the mid-late 2000s the RBNZ held the view that house prices were above their fundamental values.³⁴ Consequently, they were forecasting house prices to decrease. This is apparent in RBNZ forecasts for more than 4 quarters ahead. Figure 9 shows the RBNZ and combination forecasts graphed against actual outturns. Before the housing market downturn in 2007/08, the RBNZ’s long-run forecasts for annual house price inflation were consistently below 5 percent and often negative. However, house price inflation was persistently above 10 percent.

The strong negative bias of the RBNZ’s forecasts highlight that the RBNZ tended to place a lot of weight on the fundamental level of house prices over our sample period. Many of the models used in the combination do not have a ‘fundamental’ house price value. Hence, the combination forecasts were generally higher than the RBNZ forecasts and were consequently closer to actual house price inflation.

These results show that a combination approach can produce relatively accurate house price forecasts. However, one must be mindful that the quasi-realtime combination forecasts will inherently have an advantage over the

³⁴ Reserve Bank of New Zealand (2005) and Reserve Bank of New Zealand (2008).

Figure 9
Combination and RBNZ Forecasts



RBNZ forecasts due to the ex-post knowledge of the housing market used to select indicators and develop the model suite. Furthermore, we recognise that the evaluation period is a potential caveat to our results. In particular, the forecasts were evaluated over a 12-year period from 1998Q1 that included a prolonged upswing in the housing market and consequently our results could be biased by this sample.³⁵ Regardless of these caveats, we think our results are quite striking and suggest that a combination approach should be considered in the forecast process.

7 Conclusion

In this paper we developed a range of empirical models to forecast house price inflation. Typically, the models that produced more accurate forecasts than our AR benchmark were those that used small housing-specific data

³⁵ However, we could not generate quasi-realtime forecasts over a longer sample due to data limitations.

sets, such as the REINZ VAR, the error correction model (ECM) and the rental yield model.

We generate a combined forecast using three approaches: equal weights, MSE weights and model selection. Over our evaluation period the model selection approach had lower root mean squared forecast errors than the combination methods. This result is at odds with empirical evidence, outlined in Timmermann (2006), that simple averages often produce more accurate out-of-sample forecasts than the best-performing model in real time. However, the differences between the RMSEs of the model selection and combination forecasts were small for most horizons.

We also evaluate the combination forecast against both the individual models' and the RBNZ's forecasts of house price inflation. We find that the combination forecast is on par with the best of the individual models at most horizons, and has produced lower RMSEs than the RBNZ's forecasts. For most horizons, the forecast performance of the combination was significantly better than the forecast performance of the RBNZ.

These results suggest statistical model forecasts for house price inflation should be considered in the forecast process. Moreover, a combined forecast is beneficial because it encompasses a wide range of modelling techniques and data, while producing reasonably accurate forecasts. While the results are striking, we recognise that the evaluation period is relatively short. Furthermore, in this paper we have only focused on the combination of point forecasts. A natural extension is to consider density forecast combination. This could be done in a similar manner to McDonald and Thorsrud (2011). We leave this for future work.

References

- Abelson, P, R Joyeux, G Milunovich, and D Chung (2005), “Explaining house prices in australia: 1970-2003,” *Economic Record*, 81, S96--S103.
- Bañbura, M, D Giannone, and L Reichlin (2008), “Large bayesian vars,” *European Central Bank, Working Paper Series*, 966.
- Bates, J M and C W J Granger (1969), “The combination of forecasts,” *Operational Research Quarterly*, 20(4), 451--468.
- Bjørnland, H C, K Gerdrup, A S Jore, C Smith, and L A Thorsrud (2009), “Does forecast combination improve Norges Bank inflation forecasts?” *Norges Bank, Working Paper*, 2009/01.
- Bloor, C and T Matheson (2008), “Analysing shock transmission in a data-rich environment: A large bvar for new zealand,” *Reserve Bank of New Zealand, Reserve Bank of New Zealand Discussion Paper Series*, DP2008/09.
- Boivin, J and S Ng (2006), “Are more data always better for factor analysis?” *Journal of Econometrics*, 132(1), 169--194.
- Briggs, P and T Ng (2009), “Trends and cycles in New Zealand house prices,” *Centre for Housing Research, Aotearoa New Zealand, Paper for CHRANZ workshop*.
- Coleman, A and J Landon-Lane (2007), “Housing markets and migration in new zealand, 1962-2006,” *Reserve Bank of New Zealand, Reserve Bank of New Zealand Discussion Paper Series*, DP2007/12.
- Das, S, R Gupta, and A Kabundi (2009), “Could we have predicted the recent downturn in the south african housing market?” *Journal of Housing Economics*, 18(4), 325--335.
- De Veirman, E and A Dunstan (2008), “How do housing wealth, financial wealth and consumption interact? evidence from New Zealand,” *Reserve Bank of New Zealand, Reserve Bank of New Zealand Discussion Paper Series*, DP2008/05.
- Diebold, F X and R S Mariano (1995), “Comparing predictive accuracy,” *Journal of Business & Economic Statistics*, 13(3), 253--63.
- DTZ New Zealand (2004), “Changes in the structure of the New Zealand housing market volume 1,” *Centre for Housing Research Aotearoa New Zealand*.
- Dvornak, N and M Kohler (2007), “Housing wealth, stock market wealth and consumption: A panel analysis for Australia,” *The Economic Record*, 83(261), 117--130.
- Fraser, P, M Hoesli, and L McAlevey (2008), “House prices and bubbles in New Zealand,” *The Journal of Real Estate Finance and Economics*,

- 37(1), 71--97.
- Gupta, R, A Kabundi, and S M Miller (2009), "Using large data sets to forecast housing prices: A case study of twenty us states," *University of Pretoria, Department of Economics, Working Papers*, 200912.
- Gupta, R, A Kabundi, and S M Miller (2011), "Forecasting the us real house price index: Structural and non-structural models with and without fundamentals," *Economic Modelling*, 28(4), 2013--2021.
- Herring, R J (2006), "Booms and busts in housing markets: How vulnerable is New Zealand?" Available from: <http://www.rbnz.govt.nz/research/fellowship/3324563.pdf>.
- Hull, L (2003), "Financial deregulation and household indebtedness," *Reserve Bank of New Zealand, Reserve Bank of New Zealand Discussion Paper Series*, DP2003/01.
- Iacoviello, M and S Neri (2010), "Housing market spillovers: Evidence from an estimated dsge model," *American Economic Journal: Macroeconomics*, 2(2), 125--64.
- International Monetary Fund (2003), "United Kingdom: Selected issues," *International Monetary Fund, International Monetary Fund Country Report*, 03/47.
- LeSage, J P (1999), *Applied Econometrics using MATLAB*.
- McDonald, C and M Smith (2009), "Developing stratified housing price measures for New Zealand," *Reserve Bank of New Zealand, Reserve Bank of New Zealand Discussion Paper Series*, DP2009/07.
- McDonald, C and L Thorsrud (2011), "Evaluating density forecasts: Model combination strategies versus the RBNZ," *Reserve Bank of New Zealand, Reserve Bank of New Zealand Discussion Paper Series*, DP2011/03.
- Muellbauer, J N (2007), "Housing, credit and consumer expenditure," *Proceedings*, 267-- 334.
- O'Donovan, B and D Rae (1997), "The determinants of house prices in New Zealand: An aggregate and regional analysis," *New Zealand Economic Papers*, 31(2), 175--198.
- Reserve Bank of New Zealand (2005), "Financial Stability Report, November 2005," *Reserve Bank of New Zealand*.
- Reserve Bank of New Zealand (2008), "Monetary Policy Statement, March 2008," *Reserve Bank of New Zealand*.
- Stock, J H and M W Watson (2002), "Macroeconomic forecasting using diffusion indexes," *Journal of Business & Economic Statistics*, 20(2), 147--62.
- Timmermann, A (2006), *Forecast Combinations*, vol 1 of *Handbook of Economic Forecasting*, chap 4, 135--196, Elsevier.

Weekin, O (2004), “Asset pricing and the housing market,” *Bank of England, Quarterly Bulletin*.

Appendices

A. Data details

The table below lists all the data series used in our model suite, arranged by data source. We also list the frequency of the raw data ('m' for monthly and 'q' for quarterly) and whether the data is seasonally adjusted (denoted by s.a).³⁶ A * indicates that the series is adjusted. These adjustments are explained in the notes below the table. The series included in the factor model were not listed in section 3 so these are denoted by 'Factor'.

Quotable Value New Zealand

House Prices

1	Quarterly house price index (s.a)	q	Factor
2	Quarterly house price index - Auckland (s.a)	q	
3	Quarterly house price index -Wellington (s.a)	q	
4	Quarterly house price index - Christchurch (s.a)	q	

Statistics New Zealand

National Accounts

5	Real GDP - total production (s.a)	q	Factor
6	Nominal GDP - total production (s.a)	q	
7	Real GDP - private investment dwellings (s.a)	q	
8	Real GDP - total private consumption (s.a)	q	

Price data

9	Consumer price index	q	Factor
---	----------------------	---	--------

Migration Data

10	Permanent and Long-term migration arrivals (s.a)*	m	Factor
11	Permanent and Long-term migration arrivals (s.a)*	m	Factor

Labour market Data

12	HLFS official unemployment rate (s.a)	q	
13	Total population (s.a)	q	
14	HLFS Working age population (s.a)	q	

Housing related data

15	Number of dwellings (s.a)	q	
16	Number of consents for new dwellings (s.a)	m	

³⁶ Most models are estimated using quarterly data so for these models the monthly data is first converted to a quarterly frequency by taking a quarterly average.

17	Value of consents for new dwellings (s.a)	m	Factor
18	New dwellings - value (s.a)	m	Factor
19	New dwellings - floor area (s.a)	m	Factor
<i>Merchandise trade indexes</i>			
20	Terms of trade index	q	Factor
Reserve Bank of New Zealand			
<i>Interest and Exchange rates</i>			
21	Call rate	m	Factor
22	30 day Bank Bill yield	m	Factor
23	60 day Bank Bill yield	m	Factor
24	90-day Bank Bill yield	m	Factor
25	1 year Government Bond yield	m	Factor
26	2 year Government Bond yield	m	Factor
27	5 year Government Bond yield	m	Factor
28	10 year Government Bond yield	m	Factor
29	Floating mortgage - new customer rate	m	
30	2-year mortgage rate - new customer rate*	m	
31	5-year mortgage rate - new customer rate*	m	
32	New Zealand dollar TWI	m	Factor
<i>Survey of expectations</i>			
33	Expected annual CPI inflation - 2 years ahead	q	
Real Estate Institute of New Zealand			
<i>Housing related data</i>			
34	Stratified house price index (s.a)	m	Factor
35	Number of dwelling sales (s.a)*	m	
36	Median days to sell (s.a)	m	
Datastream			
37	Australia unemployment rate (standardised)	q	Factor
Ministry of housing			
38	Median weekly rent (s.a)*	m	
ANZ			
39	ANZ commodity price index - NZ\$	m	Factor

Series used only for data adjustment

Reserve Bank of New Zealand

Interest and Exchange rates

40	2 year swap rate	q
41	5 year swap rate	q

Statistics New Zealand

CPI data

42	Actual rentals for housing (s.a)	q
----	----------------------------------	---

Notes

- The migration series (series 10 and 11) are de-trended using the HLFS working age population (series 41).
- The 2-year and 5-year mortgage rate series (series 30 and 31) only date back to 1998M6. These series are backdated using the 2 and 5 year swap rates (series 40 and 41) plus the margin that existed between the respective mortgage and swap rates at the beginning of the sample.
- The REINZ series for the number of dwelling sales (series 34) is de-trended by subtracting an approximation of the number of new dwellings sold each month (the quarterly change in the stock of residential dwellings (series 15), divided by three to spread the sales over each month in the quarter). This adjustment leaves just sales of existing dwellings.³⁷
- The median weekly rent series (series 38) only dates back to 1993M1 and is backdated prior to this using the actual rentals for housing series from the CPI index (series 42).

³⁷ The adjusted series has a higher correlation with the REINZ house price index. We use this increase in correlation as a justification for the chosen de-trending method. We also de-trended the series using a linear trend, and the working age population but found these methods to produce inferior in-sample fits.