Reserve Bank of New Zealand
Analytical Notes

How has the LVR restriction affected the housing market: a counterfactual analysis

AN2014/03

Gael Price

May 2014

Reserve Bank of New Zealand Analytical Note series
ISSN 2230-5505

Reserve Bank of New Zealand
PO Box 2498
Wellington
NEW ZEALAND

www.rbnz.govt.nz

RESERVE BANK
OF NEW ZEALAND
TE PŪTEA MATUA

NON-TECHNICAL SUMMARY

The Reserve Bank introduced a loan-to-value ratio (LVR) ‘speed limit’ for housing loans in October 2013. It is not easy to estimate the real-world effects of a policy change. Since the LVR restriction was introduced, house sales have fallen, house price inflation has declined, and credit growth has stabilised. But other factors also affect house prices and credit growth (over this period, for example, net immigration has increased and economic activity indicators have been strong), and determining how much of the housing market changes can be attributed to LVR restrictions requires some analytical framework.

To disentangle the effects of the LVR restriction from other influences that have affected the housing market, this paper uses a counterfactual scenario – a description of what could have happened in the absence of the LVR restriction. We estimate the usual relationships between housing market variables over history, and then use those relationships to predict the likely ‘normal’ behaviour of the housing market, starting from just before the LVR policy was announced. The forecast also needs to take account of developments in the broader economy. Hence, the forecast is conditioned on other variables that are important for the housing market – migration, economic activity, and interest rates.

Actual housing market activity, house price inflation, and household credit growth have fallen below the estimated counterfactual scenario since the LVR restriction was introduced. House price inflation is 3.3 percentage points lower than the model suggests it could have been, and household credit growth is 0.9 percentage points lower (as at March 2014). This result is most likely to have occurred because of the imposition of the LVR restriction. There has been quite a rapid reaction in the housing market during the initial period of transition. It will be important to continue assessing the effects of the policy as the market adjusts.

To the extent that the model captures well the behaviour of the housing market, the results suggest that the LVR restriction has reduced house price inflation and housing-related credit growth. The framework used here relies on the assumption that the LVR restrictions have been the main source of developments within the housing market itself (as distinct from external influences such as migration or economic activity) since September 2013. This assumption will become increasingly unlikely to hold over time. To assess the effectiveness of the LVR restrictions beyond the initial adjustment period, other techniques will be needed.

1 Introduction[1]

The Reserve Bank introduced a loan-to-value ratio (LVR) restriction for housing loans from October 2013. The aim of the policy was *"to help slow the rate of housing-related credit growth and house price inflation"*, with the ultimate goal of *"reducing the risk of a substantial downward correction in house prices that would damage the financial sector and the broader economy"* (Wheeler 2013). The restriction was implemented as a 'speed limit' whereby banks can make high-LVR loans (more than 80 percent of the house value) up to 10 percent of the value of their new mortgage lending over any three-month period (Rogers 2014).

Enough time has now passed to allow an assessment of the impact of the policy. Credit growth and house price inflation have slowed in the months since the LVR restriction was introduced late last year (figure 1). However, other factors also affect the housing market, making it harder to determine the effects of the LVR restriction itself. We approach the problem by estimating a counterfactual scenario – a description of what might have happened to key housing market variables between September 2013 and March 2014, taking account of some of the other key external factors that influence the housing market, if the LVR restriction had not been introduced.

2 Scope of this paper

It is not easy to answer empirical questions about the actual effects of a policy change. Although we may see changes in the behaviour of the data once the policy is introduced, it is very difficult to tell whether that behaviour is caused by the policy, or would have occurred anyway. This problem is often expressed as 'we don't know the counterfactual' – that is, we do not know what would have happened had the policy not been introduced. Without knowing that, it is impossible to say just what the effects of the policy have been.

[1] I would like to thank Anella Munro, Chris Bloor, Michael Reddell, Michelle Lewis, Miles Parker, Nick Sander, Ross Kendall, Yuong Ha, and Richard Sullivan for comments and suggestions.

Figure 1: House price inflation and household credit growth

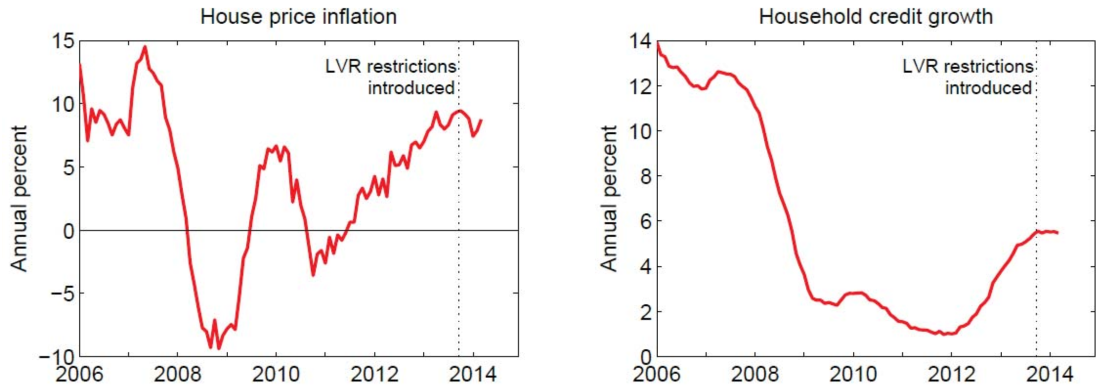

Source: RBNZ, Real Estate Institute of New Zealand (REINZ).

The true counterfactual is unobservable. To know it with certainty would require a perfect model of the economy – an impossibility. However, there is observable information that can help us to estimate a counterfactual scenario (with some uncertainty). The behaviour of relevant variables before the policy was announced, and developments in the broader economy since the policy was announced, both help us to get at the counterfactual. Using the past relationships between variables of interest (housing market variables), we can forecast how those variables could have evolved if their normal behaviour had prevailed without being influenced by the LVR policy. Using information about recent developments in the broader economy, we can let the forecast be influenced by other important data that affect the housing market.

In this way, the counterfactual scenario gives an estimate of what could have happened in the housing market if the policy had not been introduced, but the wider economy had still evolved as described by the conditioning information.

The conditioning information needs to include variables that can influence the housing market, while not themselves being influenced by housing market developments – ‘exogenous’ variables. To estimate the effect of a policy, it will also be important that the conditioning variables are not affected by the policy itself. The key exogenous variables for the housing market are migration flows, economic activity, and interest rates (McDonald 2013). Since LVR restrictions were introduced, migration and activity have been very strong, providing a positive impetus to the housing market (figure 2). However, interest rates have also begun to rise, possibly offsetting some of the boost from the other exogenous factors. The counterfactual scenario is conditioned on all these exogenous variables in order to estimate the net effect on the housing market.

By conditioning the counterfactual scenario on data of this sort, we can allow the estimated counterfactual to respond to the fundamental developments captured by those data series – things like migrants' desire to live in New Zealand, which affects housing demand; firms' beliefs about the future, which affect economic activity and hence homebuyers' incomes; and the level of spare capacity in the economy, which affects interest rates and hence the cost of borrowing to buy a home. We assume that, over the sorts of horizons discussed here, the conditioning variables themselves will not have been materially influenced by the LVR restrictions.[2]

Figure 2: Data used to condition the counterfactual scenario[3]

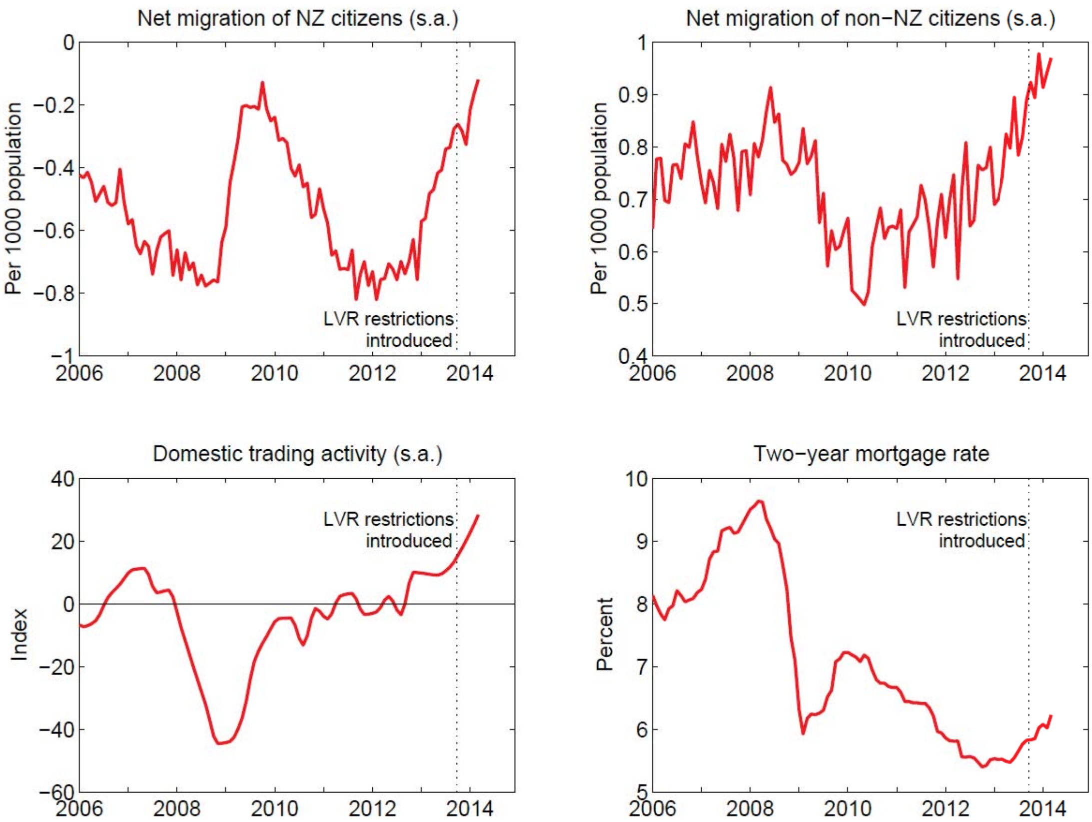

Source: New Zealand Institute of Economic Research Ltd (NZIER), RBNZ, Statistics New Zealand.

[2] Assuming that business confidence and migration have not been materially affected by LVR restrictions seems reasonable. The assumption that interest rates are unaffected by LVR restrictions is more questionable. In reality it is likely that interest rates will be lower in the near term than they would have been without LVR restrictions, as a direct result of the policy (Spencer 2014). If this effect has not yet appeared in the data, there are no implications for the current analysis. However, if two-year mortgage interest rates are already lower than they would have been in the absence of LVR restrictions, that would mean the counterfactual scenario should be conditioned on higher interest rates than those we actually observe in order to accurately capture a world without LVR restrictions. To test the sensitivity of the results to this factor, the analysis was repeated with interest rates 50 basis points higher than actually observed from October 2013 to March 2014. The results are almost unchanged (see appendix B).

[3] See appendix A for more detailed descriptions of the data.

By conditioning the counterfactual scenario on the exogenous data, we allow for important developments that originate outside the housing market. However, we are not allowing for developments that originate inside the housing market. There is no source of information about these intrinsic housing market ‘shocks’ that could be occurring, such as housing supply shocks. Consequently, the counterfactual scenario is constructed under the implicit assumption that such intrinsic shocks have not occurred over the relatively short period in question.

If this assumption is incorrect, the gap between the data and the counterfactual scenario will represent the combined effect of the LVR restrictions and the intrinsic shocks. It is difficult to disentangle the different causes from one another, and so, in the language often used, the policy’s effect is unidentified – it cannot be distinguished from the effects of intrinsic shocks.

This paper does not attempt to identify the different causes of deviations between the data and the counterfactual scenario. The shocks induced by the LVR restrictions are entirely new to the New Zealand economy, so there is little information in past data to guide any attempt at identification. Instead, this paper makes the assumption that LVR restrictions have been the major source of shocks inside the housing market itself over the past six months.

There are two reasons why this is likely to be a reasonable assumption. Firstly, there have been few other policy interventions or unusual housing market developments since LVR restrictions were introduced. The Auckland Housing Accord does not appear to have a significant effect on actual housing supply as yet, and although rebuild work in Canterbury could also complicate the picture, results estimated excluding Canterbury are very similar to whole-of-New Zealand results. Consequently, LVR restrictions are the largest known unusual influence currently acting inside the housing market.

The second reason for making this assumption stems from the nature of the empirical framework we are using. By basing the analysis on estimated historical relationships prior to the announcement of LVR restrictions, we are assuming that these historical relationships still hold. In reality, the policy is likely to have changed the relationships between variables. As the data evolve according to the new relationships, there are likely to be large deviations from what would have been predicted by the old relationships. Deviations of this type are generally larger than those arising from short-lived shocks. As a result, the largest deviations are likely to be those arising from the policy change and its associated change in relationships.

The framework used in this paper is unlikely to remain useful very far beyond the near term. As time goes on, two things will happen to undermine the assumptions on which the analysis relies. Firstly, intrinsic housing market shocks will occur in each period, and eventually the sum of their effects will be larger than the effect of the policy. At that stage the assumption that LVR restrictions are the major source of shocks will no longer hold. Secondly, the historical relationships being used to condition the forecast will become less relevant as the economy transitions to new relationships that reflect the new policy environment. As there is not enough post-implementation data to estimate these new relationships, we are forced to keep assuming the old relationships, with the result that the validity of the conditioned forecast will decline over time.

The current juncture – six months after the LVR restriction was first implemented – is a convenient time to assess the effects of the policy. Any sooner, and there would have been insufficient data to draw clear conclusions; much later, and the limitations of our framework will make the results less likely to be valid. Further assessment of the effect of the LVR restrictions beyond the near term will require proper identification of the various shocks affecting the housing market.

3 Approach

This paper uses statistical techniques to estimate the historical relationships between variables of interest and to apply those relationships over the period since LVR restrictions were introduced. We use a vector autoregression, or VAR – a time series framework in which every variable is modelled as a function of its own lags and lags of the other variables.

The model is similar to that used in the Bloor and McDonald (2013) initial assessment of the likely effects of LVR restrictions, but uses the dataset of McDonald (2013) wherever possible, as that more recent model was explicitly focused on understanding developments in the housing market.

We estimate the model on seasonally adjusted monthly data from January 1992 to April 2013. We include six lags of each variable, and use Bayesian techniques to estimate the parameters.[4] The end date is chosen to exclude the period after the RBNZ's macroprudential policy framework was finalised in May 2013 (RBNZ 2013), in case the

[4] The lag length and Bayesian estimation approach is the same as Bloor and McDonald (2013), with random walk priors implemented using dummy variables.

discussion of LVR restrictions in the policy document led market participants to change their behaviour.

The variables used are[5,6]:

- net migration of New Zealand citizens;
- net migration of non-New Zealand citizens;
- net experienced domestic trading activity[7];
- two-year fixed mortgage rate;
- volume of house sales;
- median days to sell a house;
- residential consent issuance;
- REINZ stratified house price index; and
- household credit.

The counterfactual scenario is constructed as a forecast starting in September 2013. This start date is chosen to allow us to estimate any pre-implementation effects that may have occurred in the housing market between the announcement of LVR restrictions on 20 August and implementation on 1 October. The forecast is conditioned on all the available data for four ('exogenous') variables: the two migration series, domestic trading activity, and mortgage rates. These variables capture the major influences on the housing market that do not originate in the housing market itself (McDonald 2013). The conditioning employs the widely used Waggoner and Zha (1999) algorithm for conditional forecasting.

By conditioning on the four exogenous variables, we obtain an estimate of the counterfactual for the other five ('endogenous') variables in the model: house sales, days to sell, consent issuance, house prices, and household credit. As well, we also construct a counterfactual scenario for two off-model ('satellite') variables – housing-related credit (as opposed to total household credit, which is in the model), and the number of housing loan (mortgage)

[5] Plots and definitions of the data are in Appendix A. The two migration series are both defined as permanent and long-term (PLT) migration. The two migration series, the volume of house sales, and residential consent issuance, are all expressed per 1000 population.

[6] To test the sensitivity of the results to the particular data used, the analysis was repeated with a selection of alternative series. In all cases the results were almost unchanged (see appendix B).

[7] From the NZIER Quarterly Survey of Business Opinion (QSBO). McDonald (2013) uses the output gap as the measure of capacity pressures, but for the current exercise it is important to have a stable real-time measure, as we are trying to assess the effect of a policy in near-real time. As output gap estimates are subject to large revisions, we substitute the QSBO series, which is seldom revised.

approvals per 1000 population. These variables are both of interest when assessing the effects of the LVR policy, but their time series are too short to be included in the model. Instead, we model each satellite variable separately as a function of its own lags and all the in-model variables. The counterfactual scenario for these variables is a forecast conditioned on the counterfactual scenario for all the in-model variables, as well as on the conditioning data.

The results in the next section are presented with error bands (dashed lines) around the counterfactual scenario for each variable. These bands represent the 95 percent confidence interval for possible counterfactual scenarios, based on the historical forecast errors of the model. If the data lies outside the bands around the counterfactual scenario, that means the data is so extreme that the model did not predict such an outcome in 95 percent of possible forecasts, given its usual ability to predict the data. If the data lies inside the bands, its deviation from the counterfactual scenario is comparable to the model's usual forecast error, and so the data is not distinguishable from the counterfactual scenario in a statistical sense.

The approach used for drawing error bands means that they are wide and asymmetrical.[8] These characteristics arise from the short sample period used. Over the sample period, the model's forecast errors do not have a mean of zero – if, for example, the bands are skewed upward, there were more upside than downside errors over the sample period. In addition, over the short sample period, episodes of volatility make up a higher proportion of the data than they would in a long sample, where volatile periods would be balanced by stable periods. Consequently, the uncertainty inherent in the data is large in the shorter sample, generating wide error bands. The wide bands mean that the threshold for statistical significance is high – it is difficult to get a significant result using this short sample, even if the result 'should' be significant in a larger sample. Thus, the significant results that are obtained are reasonably reliable.

4 Results

The housing market began to weaken relative to the counterfactual scenario as soon as LVR restrictions were implemented (figure 3 and table 1). Activity declined, with house sales and mortgage approvals falling significantly below the counterfactual scenario, and days to sell increasing. House price inflation declined below the counterfactual scenario, although not significantly – as at March 2014, annual house price inflation was 3.3 percent below the counterfactual scenario. Annual household credit growth fell modestly below the

[8] See appendix C for a full description of the procedure used to estimate the error bands.

counterfactual scenario, reaching 0.9 percent below by March 2014, while housing-related credit growth also declined below the counterfactual scenario. Residential building consents rose above the counterfactual scenario.

The fall in house sales was larger than estimated in Bloor and McDonald (2013). The large fall may have arisen from a rapid reaction to the policy by market participants, as has been seen with the banks (in aggregate) complying with the high-LVR ‘speed limit’ several months before they were required to. It is possible that house sales may recover as the market adjusts to the new policy regime, especially if banks make more use of exemptions, as was initially expected (Rogers 2014).

House prices have not yet declined significantly below the counterfactual scenario. One possibility is that this may partly be explained by an upward bias in the measurement of house prices. The best house price measure in New Zealand is the Quotable Value New Zealand (QVNZ) measure, which is not affected by the composition of houses sold in any period, but is reported with a significant lag. Thus in this exercise we use the more timely stratified REINZ measure. This measure corrects for compositional changes in sales from district to district, but not for changes in the composition of houses sold within any particular district. Indications are that the slowdown in house sales in recent months has been concentrated in lower-priced houses (e.g. those typically purchased by first home buyers). If so, this will be biasing the REINZ measure upwards relative to the QVNZ measure. That would imply that the true effect of LVR restrictions on house prices could be larger than estimated here, although probably not by much.[9]

Household credit growth has responded quite slowly to the weaker housing market, partly because there is a lag of several months between the sale of a house and the corresponding housing loan being drawn down. Credit growth is likely to fall further below the counterfactual scenario in coming months, based on the decline in activity already seen (see figure 4).

The sharp rise in residential building consents relative to the counterfactual scenario, although not statistically significant, suggests an alternative interpretation for the results. As there has been a rise in consents and a decline in prices (relative to the counterfactual scenario), this suggests a housing supply shock may be at work. However, increased

[9] Between the September and December quarters of 2013, house prices rose by 2.4 percent on the REINZ measure, but only 1.2 percent on the QVNZ measure. A deviation of this size between the two measures is not very unusual, having typically occurred about once every two years over the history of the series.

housing supply would not lead to the observed declines in house sales or credit growth – those results are best explained by the LVR restrictions.

Table 1: Deviations of housing market data from the counterfactual scenario

	House sales per 1000 population (%)	Number of mortgage approvals per 1000 population (%)	Days to sell (days)	Residential building consents per 1000 population (%)	House price inflation (annual percentage points)	Household credit growth (annual percentage points)	Housing-related credit growth (annual percentage points)
Sep 2013	0.7	-4.9	-0.9	9.0	-0.2	0.1	0.0
Oct 2013	-10.3	0.1	0.6	12.4	0.1	0.1	0.0
Nov 2013	-16.7 (*)	-12.3 (*)	1.5	21.4	-0.5	0.1	0.0
Dec 2013	-21.2 (*)	-8.8 (*)	3.0	35.5	-1.7	-0.1	-0.3
Jan 2014	-22.7 (*)	-7.9 (*)	4.6	12.4	-3.3	-0.3	-0.5
Feb 2014	-25.8 (*)	-10.4 (*)	5.1	7.7	-3.8	-0.5	-0.6
Mar 2014	-26.2 (*)	-12.1 (*)	7.0	3.0	-3.3	-0.9	-1.0

Note: numbers in red indicate that the data is below the counterfactual scenario, while numbers in black indicate that the data is above the counterfactual scenario. Numbers marked with an asterisk (*) indicate that the data is significantly different from the counterfactual scenario at the 95 percent confidence level.

Two of the endogenous variables tend to respond with a lag to changes in housing market activity – house price inflation and credit growth. The model's forecast for these variables, conditioned on all the available data (not the limited dataset used in figure 3), suggests that they are both likely to decline further below the counterfactual scenario in the near future as a result of the declines in activity that have already occurred (figure 4). A forecast of this type expresses the expected near-term impacts of recent developments, but is based on the assumptions that historical relationships continue to hold, and that no further shocks occur.[10]

[10] Beyond the lag horizon of the model – six months in our case – the forecast will show the variables returning to their historical average behaviour. This is unhelpful when we know that the variables are in reality being affected by LVR restrictions, meaning that historical average behaviour is unlikely to be relevant. Consequently, the forecasts in figure 4 end at the six-month horizon.

Figure 3: Counterfactual scenario vs actual data (endogenous and satellite variables)

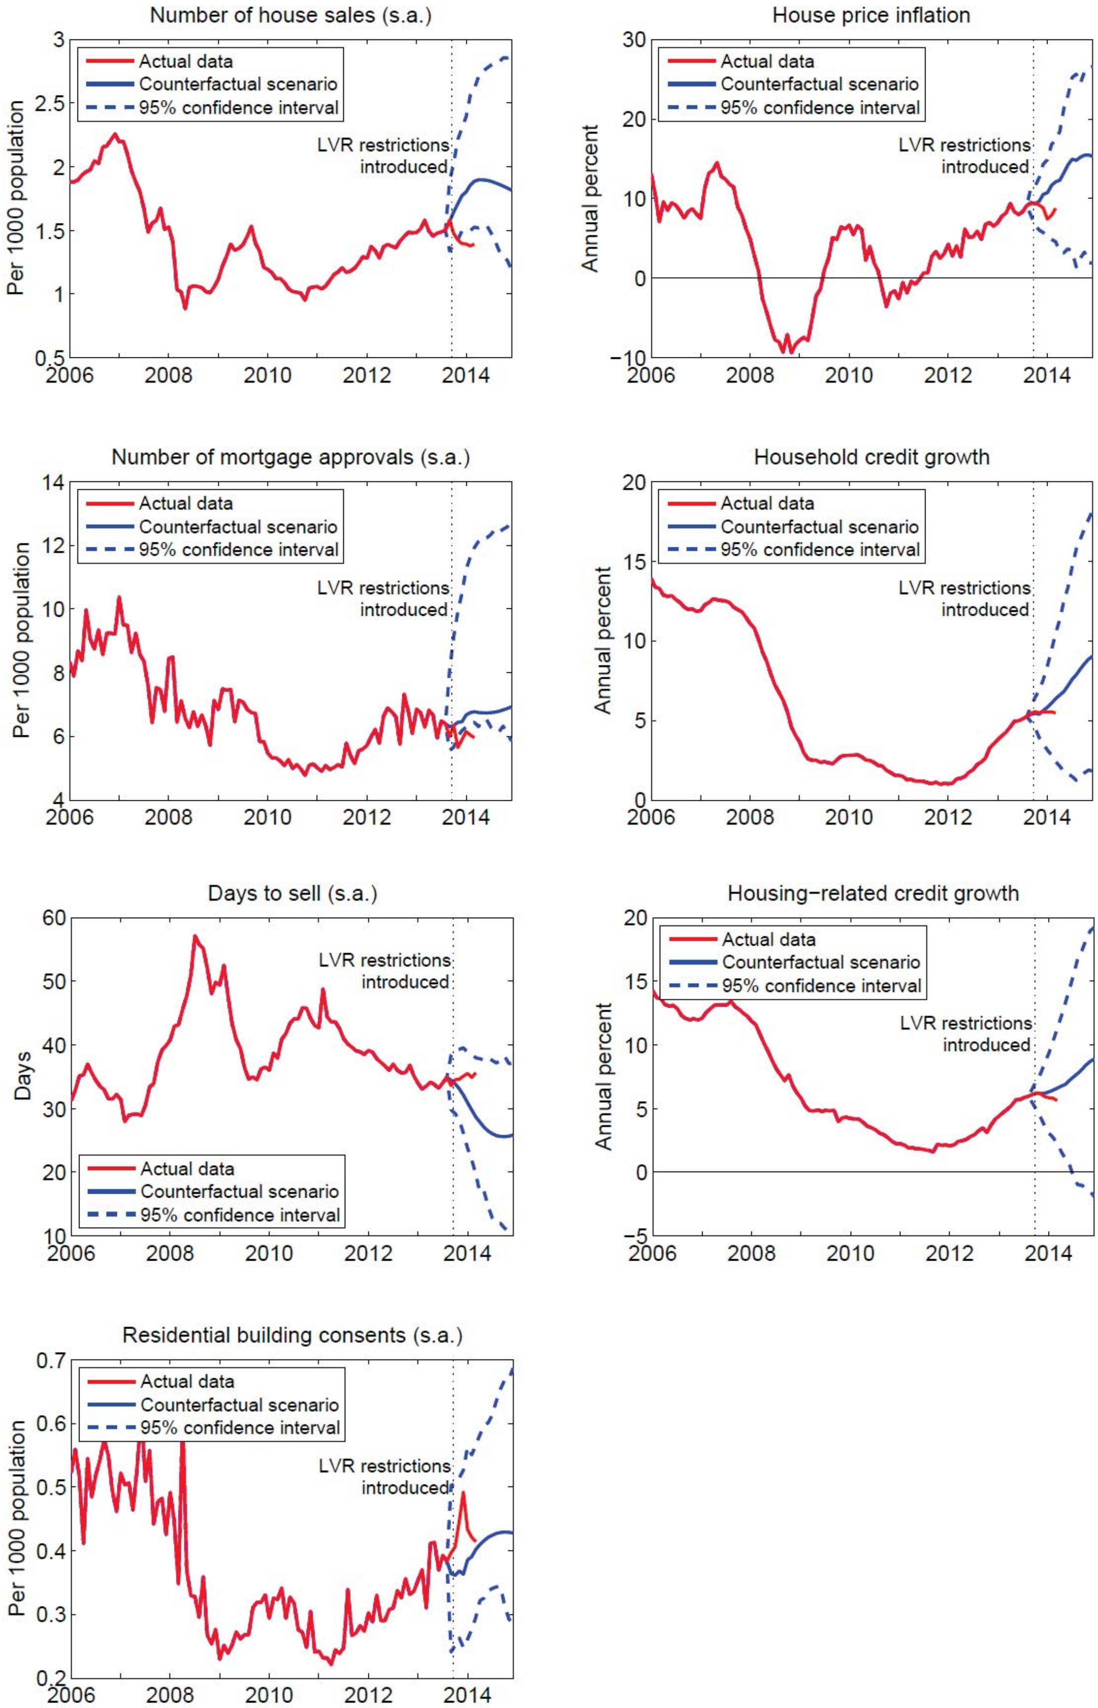

Source: RBNZ, REINZ, Statistics New Zealand, author’s estimates. Dashed lines show the range with a 95 percent change of containing the true counterfactual, based on historical forecast errors.

Figure 4: Counterfactual scenario vs forecast (house price inflation and household credit growth)

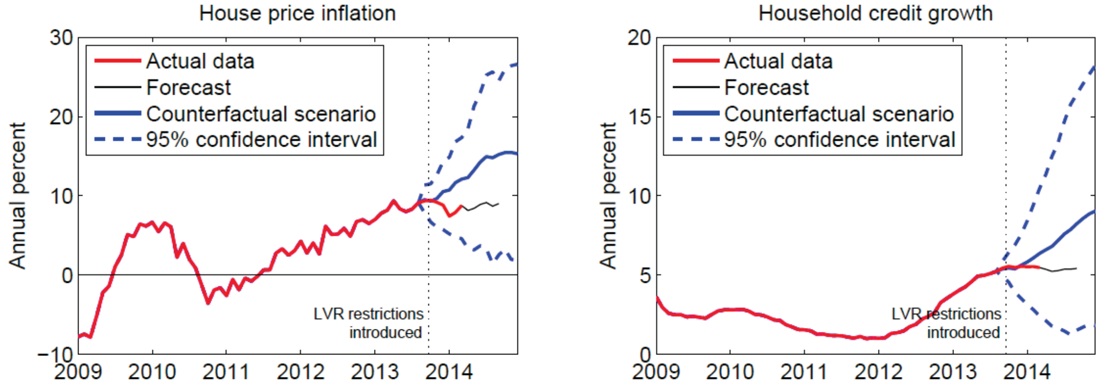

Source: RBNZ, REINZ, author's estimates. Dashed lines show the range with a 95 percent change of containing the true counterfactual, based on historical forecast errors. The forecast is a simple statistical model-based forecast conditioned on all the data up to March 2014.

5 CONCLUSION

The LVR 'speed limit' was implemented with the aim of reducing house price inflation and household credit growth, and thereby improving financial system resilience. Now that six months have elapsed since the introduction of the LVR restriction, it appears that the policy has moderated house price inflation and credit growth. An estimated counterfactual scenario suggests that, in the absence of the LVR policy or any other housing-specific shocks, house price inflation could have been 3.3 percentage points higher and household credit growth could have been 0.9 percentage points higher (on an annual basis to March 2014). The LVR restriction is the most likely explanation for that result.

The estimated effects in this paper are in line with those from Bloor and McDonald's (2013) initial estimates of what the effects of LVR restriction might be. That analysis showed that, over the first year, house price inflation was likely to fall by 1-4 percentage points, and credit growth by 1-3 percentage points. As both variables may decline further in the current analysis, it is possible that the effect on house price inflation by the end of the first year may be somewhat larger than originally predicted. However, the first six months after implementation probably represent a transitional period, during which market participants may have reacted quite rapidly to the policy. As participants grow accustomed to the new policy environment, it is possible that housing market activity could rebound somewhat, leading to smaller effects over the first year than those implied by the analysis in this paper.

The framework of this paper is one way of assessing the impact of the LVR restriction over the initial adjustment period. The analysis relies on the assumption that the LVR restriction is likely to been the biggest source of shocks within the housing market itself (as distinct from non-housing sector shocks such as changes in net immigration) over that period. In order to assess the effects of the LVR policy beyond the near term, it will be necessary to explicitly separate LVR-related developments from other influences on the housing market.

REFERENCES

Bloor, C. and C. McDonald (2013) 'Estimating the impacts of restrictions on high LVR lending', Reserve Bank of New Zealand Analytical Note, 2013/05.

McDonald, C. (2013) 'Migration and the housing market', Reserve Bank of New Zealand Analytical Note, 2013/10.

RBNZ (2013) 'A new macro-prudential policy framework for New Zealand – final policy position', http://www.rbnz.govt.nz/financial_stability/macro-prudential_policy/5270425.pdf, May 2013.

Rogers, L. (2014) 'An A to Z of loan-to-value ratio (LVR) restrictions', Reserve Bank of New Zealand *Bulletin*, 77(1).

Spencer, G. (2014) 'Coordination of monetary policy and macro-prudential policy', http://www.rbnz.govt.nz/research_and_publications/speeches/2014/5671407.html, speech delivered to Credit Suisse Asian Investment Conference in Hong Kong, 27 March 2014.

Waggoner, D.F. and T. Zha (1999), 'Conditional forecasts in dynamic multivariate models', *The Review of Economics and Statistics* 81(4), 639-651.

Wheeler, G. (2013) 'The introduction of macro-prudential policy', http://www.rbnz.govt.nz/research_and_publications/speeches/2013/5407267.html, speech delivered to Otago University in Dunedin, 20 August 2013.

APPENDIX A - DATA

This paper uses the following variables:

- Conditioning data (figure A1)
  - net PLT migration of New Zealand citizens per 1000 population;
  - net PLT migration of non-New Zealand citizens per 1000 population;
  - net experienced domestic trading activity (QSBO);
  - two-year fixed mortgage rate;

- Endogenous variables (figure A2 and A3)
  - number of house sales per 1000 population;
  - median days to sell a house;
  - number of residential building consents issued (new buildings);
  - REINZ stratified house price index;
  - household credit;

- Satellite variables (figure A4)
  - number of mortgage approvals per 1000 population; and
  - housing-related credit.

House prices, household credit, and housing-related credit enter the model in log levels, and results for these variables have been transformed to annual growth rates for presentation purposes. All other variables enter the model in levels. All series except the mortgage rate have been seasonally adjusted. Domestic trading activity has been interpolated from quarterly to monthly frequency.

Household credit and housing-related credit are both RBNZ series from the Registered Bank Standard Statistical Return (SSR). Housing-related credit measures loans made for housing purposes as reported by banks, and makes up 85 to 95 percent of the total household credit series.

The number of mortgage approvals is an experimental RBNZ series, available from http://www.rbnz.govt.nz/statistics/ in table C16. It measures housing loans approved by banks, but not necessarily drawn on, and includes top-up loans as well as second and lower-ranking housing loans. Consequently, this series is typically around five times larger than the number of house sales.

Figure A1: Conditioning data

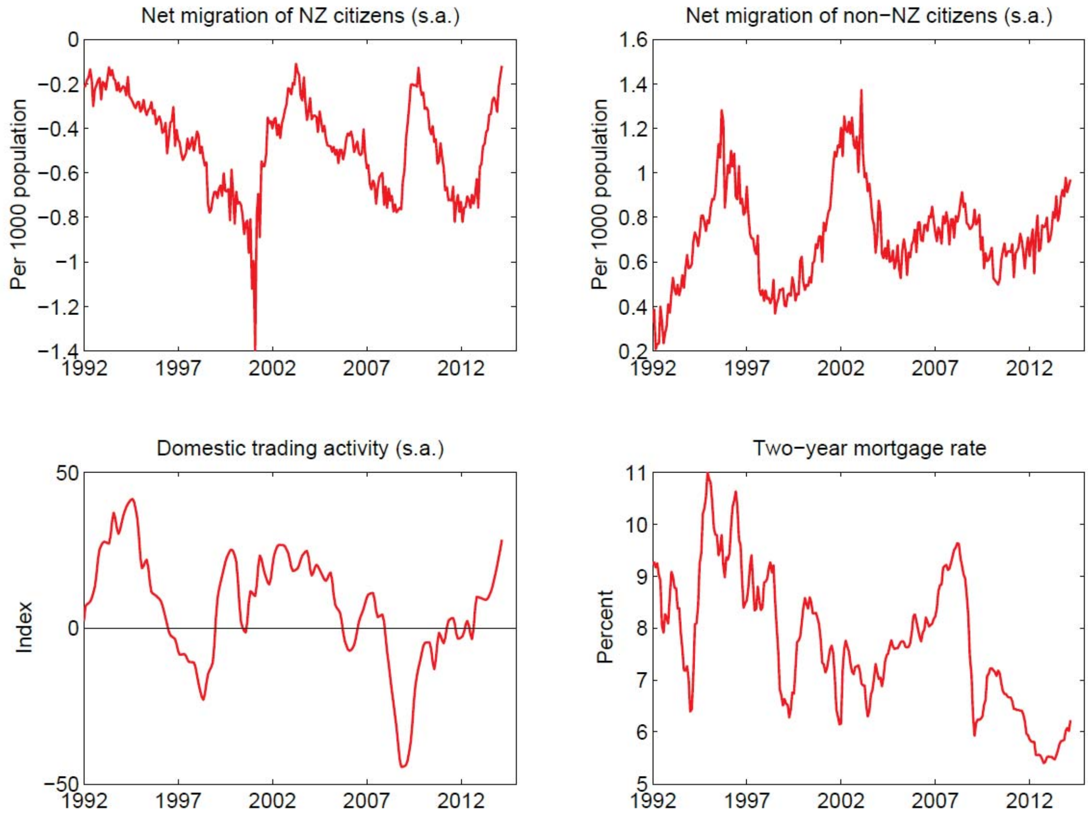

Source: NZIER, RBNZ, Statistics New Zealand.

Figure A2: Endogenous variables in levels

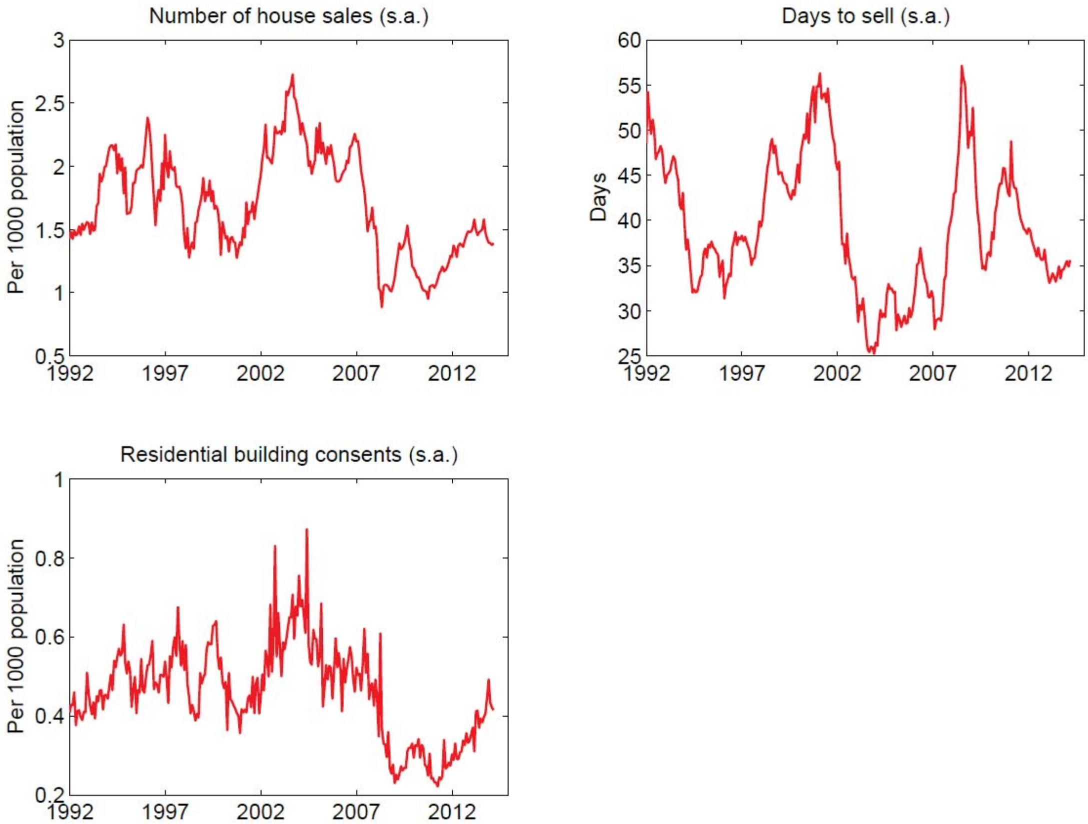

Source: REINZ, Statistics New Zealand.

Figure A3: Endogenous variables in log levels

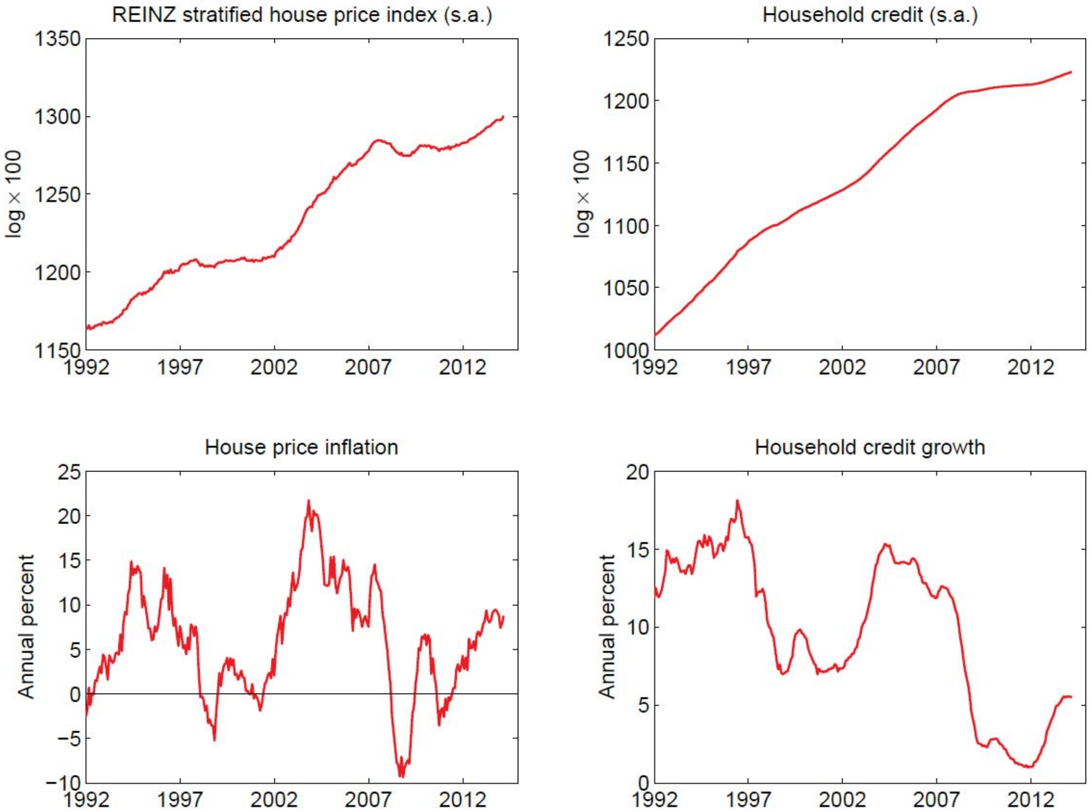

Source: RBNZ, REINZ.

Figure A4: Satellite variables

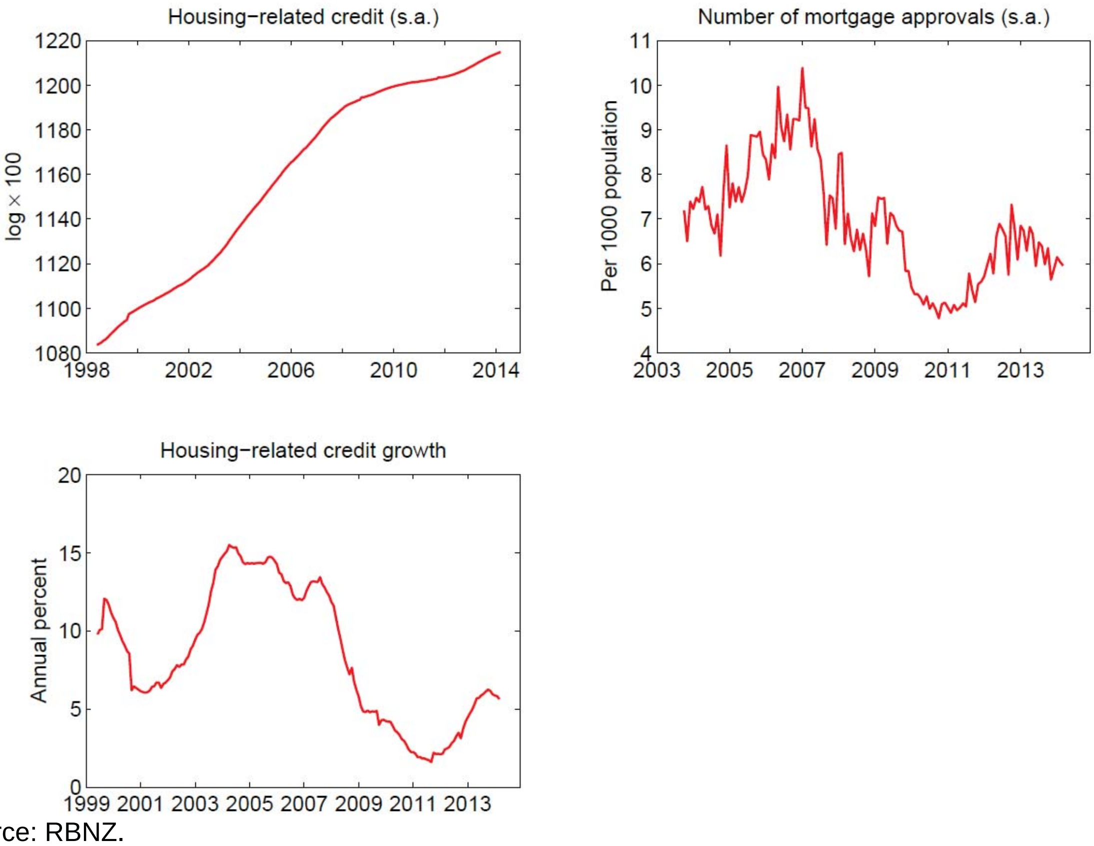

Source: RBNZ.

APPENDIX B – ROBUSTNESS CHECKS

To test the sensitivity to various assumptions, the analysis was repeated five times:

1. Using an arrivals/departures split of net PLT migration, instead of the citizens/non-citizens split, following another model from McDonald (2013) (table B1).
2. Using a residential consents series that excluded consents for the Canterbury region, to ensure the analysis is not distorted by rebuild work in Canterbury (table B2).
3. Using a residential consents series that excluded apartments, to ensure the analysis is not distorted by the volatile apartments series (table B3).
4. Using an interest rate series that is adjusted to be 50 basis points higher than actually observed from October 2013 to March 2014, to capture simplistically the possibility that LVR restrictions have reduced actual interest rates (table B4).
5. Using a start date of June 2013 instead of September 2013, to estimate any pre-implementation effects arising from the publication of the policy document in May (RBNZ 2013) (table B5).

In all cases the key results were almost unchanged from those presented in section 4.[11]

Table B1: Deviations of housing market data from the counterfactual scenario: arrivals/departures split of migration

	House sales per 1000 population (%)	Number of mortgage approvals per 1000 population (%)	Days to sell (days)	Residential building consents per 1000 population (%)	House price inflation (annual percentage points)	Household credit growth (annual percentage points)	Housing-related credit growth (annual percentage points)
Sep 2013	1.9	-3.3	-0.9	6.4	-0.3	0.1	0.0
Oct 2013	-9.3	1.3	0.6	12.3	0.1	0.1	0.0
Nov 2013	-15.5	-10.1 (*)	1.5	22.5	-0.5	0.1	0.0
Dec 2013	-19.5 (*)	-6.9	2.9	35.0	-1.8	-0.1	-0.2
Jan 2014	-21.6 (*)	-2.9	4.3	14.2	-3.4	-0.3	-0.4
Feb 2014	-24.6 (*)	-5.1	4.8	9.9	-3.8	-0.5	-0.6
Mar 2014	-25.1 (*)	-6.3	6.7	5.4	-3.3	-0.9	-0.9

Table B2: Deviations of housing market data from the counterfactual scenario: excluding consents for Canterbury

	House sales per 1000 population (%)	Number of mortgage approvals per 1000 population (%)	Days to sell (days)	Residential building consents per 1000 population (%)	House price inflation (annual percentage points)	Household credit growth (annual percentage points)	Housing-related credit growth (annual percentage points)
Sep 2013	0.8	-5.0	-0.8	0.3	-0.2	0.1	0.0
Oct 2013	-10.0	0.3	0.9	2.2	0.0	0.1	0.1
Nov 2013	-16.3	-12.1 (*)	1.8	18.0	-0.6	0.1	0.0
Dec 2013	-20.8 (*)	-8.3 (*)	3.3	20.1	-1.9	-0.1	-0.2
Jan 2014	-22.2 (*)	-7.5 (*)	5.0	2.1	-3.5	-0.3	-0.4
Feb 2014	-25.3 (*)	-10.0 (*)	5.5	-2.5	-4.0	-0.5	-0.6
Mar 2014	-25.7 (*)	-11.5 (*)	7.3	-9.8	-3.6	-0.9	-0.9

[11] In all tables, numbers in red indicate that the data is below the counterfactual scenario, while numbers in black indicate that the data is above the counterfactual scenario. Numbers marked with an asterisk (*) indicate that the data is significantly different from the counterfactual scenario at the 95 percent confidence level.

Table B3: Deviations of housing market data from the counterfactual scenario: excluding consents for apartments

	House sales per 1000 population (%)	Number of mortgage approvals per 1000 population (%)	Days to sell (days)	Residential building consents per 1000 population (%)	House price inflation (annual percentage points)	Household credit growth (annual percentage points)	Housing-related credit growth (annual percentage points)
Sep 2013	0.9	-4.4	-1.0	5.0	-0.1	0.0	0.0
Oct 2013	-9.8	1.2	0.3	4.0	0.3	0.1	0.0
Nov 2013	-16.3	-11.4 (*)	1.0	3.5	-0.2	0.0	-0.1
Dec 2013	-20.8 (*)	-8.0 (*)	2.4	13.9	-1.3	-0.1	-0.3
Jan 2014	-22.3 (*)	-6.2 (*)	3.9	9.8	-2.8	-0.3	-0.5
Feb 2014	-25.4 (*)	-8.8 (*)	4.3	11.2	-3.2	-0.6	-0.7
Mar 2014	-25.9 (*)	-10.1 (*)	6.1	9.9	-2.7	-0.9	-1.0

Table B4: Deviations of housing market data from the counterfactual scenario: using higher interest rates Oct 2013 to Mar 2014

	House sales per 1000 population (%)	Number of mortgage approvals per 1000 population (%)	Days to sell (days)	Residential building consents per 1000 population (%)	House price inflation (annual percentage points)	Household credit growth (annual percentage points)	Housing-related credit growth (annual percentage points)
Sep 2013	1.4	-4.2	-0.9	9.5	-0.1	0.1	0.0
Oct 2013	-9.1	1.3	0.5	13.2	0.3	0.1	0.1
Nov 2013	-15.2	-11.2 (*)	1.3	22.7	-0.3	0.1	0.0
Dec 2013	-19.7 (*)	-7.7 (*)	2.6	37.5	-1.4	-0.1	-0.2
Jan 2014	-21.0 (*)	-6.9 (*)	4.1	14.2	-2.9	-0.4	-0.4
Feb 2014	-24.0 (*)	-9.7 (*)	4.4	9.8	-3.3	-0.6	-0.6
Mar 2014	-24.3 (*)	-11.3 (*)	6.1	5.1	-2.7	-1.0	-1.0

Table B5: Deviations of housing market data from the counterfactual scenario: scenario starting in June 2013

	House sales per 1000 population (%)	Number of mortgage approvals per 1000 population (%)	Days to sell (days)	Residential building consents per 1000 population (%)	House price inflation (annual percentage points)	Household credit growth (annual percentage points)	Housing-related credit growth (annual percentage points)
Jun 2013	-2.4	-8.1	-0.4	0.3	-0.3	0.0	0.0
Jul 2013	-3.9	-2.2	0.2	4.8	-0.6	0.0	0.0
Aug 2013	-6.7	-3.6	1.3	4.8	-0.1	0.0	0.0
Sep 2013	-5.0	-8.8	0.5	9.4	-0.4	0.0	0.0
Oct 2013	-14.7	-3.8	2.2	11.5	-0.3	0.0	0.0
Nov 2013	-20.3 (*)	-15.8 (*)	3.1	19.1	-1.1	0.0	-0.2
Dec 2013	-24.2 (*)	-12.4 (*)	4.6	32.0	-2.4	-0.3	-0.4
Jan 2014	-25.2 (*)	-11.2 (*)	6.2	9.2	-4.1	-0.5	-0.7
Feb 2014	-28.0 (*)	-13.6 (*)	6.7	4.3	-4.6	-0.8	-0.9
Mar 2014	-28.2 (*)	-15.0 (*)	8.5	-0.3	-4.2	-1.2	-1.3

APPENDIX C – ERROR BANDS

The results presented in section 4 include error bands drawn from the historical forecast errors of the model. The forecast errors are obtained as follows.

1. For each month in history, from the 60th to the last:
    a. Trim the dataset back to that month.
    b. Re-estimate the model on the shortened dataset.
    c. Construct forecasts, conditioned on the next seven months of data for the conditioning variables.
    d. Calculate the forecast errors.
    e. If the estimated model is stable, save the forecast errors by horizon. If it is unstable, and hence the forecasts are likely to be explosive, discard the errors from this month.
2. For each variable:
    a. Draw 10,000 errors at each horizon from the saved historical forecast errors for that horizon.
    b. Add the errors to the central counterfactual scenario for that variable to create a range of 10,000 scenarios.
    c. Calculate the 2.5th and 97.5th percentile of the range of scenarios.

The error bands thus represent the historical forecasting performance of the model, conditioned in the same way as the central counterfactual scenario. Unstable models were estimated in 20 percent of the historical sample, so the forecasts from those periods were discarded. Conditioning the forecasts over a short sample period tends to create asymmetrical errors, so the bands are asymmetrical.

Some variables have very wide error bands. These are an artefact of estimating the model over short samples throughout history. Large forecast errors arise from specific periods in history when those variables rose or fell sharply, but over a longer sample such episodes would normally be offset by periods when the data is relatively stable, generating smaller forecast errors. The wide bands simply imply that the model does not have enough historical data to determine that the extreme events shown are unlikely. They do not imply that the data could actually reach the extremes depicted by the bands.