A Structured Approach to Stress Testing Residential Mortgage Portfolios

Tim Hampton and Ian Harrison
Financial Stability Department
Reserve Bank of New Zealand

Prepared for Reserve Bank of New Zealand workshop:
Financial Sector Balance Sheets and Vulnerability to Financial Crises

25 September 2006
Wellington, New Zealand

1. Introduction

Stress testing has been used for many years to help firms better understand potential vulnerabilities associated with their business. It is also becoming an increasingly important component of financial stability assessment frameworks. The most common and widely understood stress testing is the market risk analysis associated with holding financial assets. The same principles apply to other forms of risk, including operational and credit risk, ie estimating the potential effects of various scenarios on a range of performance indicators (eg profit/loss).

A recent New Zealand example was the stress testing of the banking system as part of the International Monetary Fund's (IMF) Financial System Assessment Programme (FSAP) in 2003. FSAPs have been performed by most IMF member countries, and are intended to help countries identify and remedy structural weaknesses in the financial sector, and thereby enhance their resilience to macroeconomic shocks and cross-border contagion. As part of the New Zealand FSAP, the Reserve Bank of New Zealand (RBNZ) provided New Zealand's systemically important banks a number of macroeconomic scenarios and worked with the banks to assess the impact of the scenarios on their balance sheets and the banking system as a whole.[1]

This exercise proved very informative for both the RBNZ and banks that participated, and provided several lessons. One in particular was the fact that the banks calculated the impact of the shock on their own performance independently meant that different approaches were used and it was difficult to identify whether different expected losses between banks were a result of different lending practices, different methodologies, or different assumptions and judgements. Furthermore, relying on 'expert judgement' means that it is very resource intensive to replicate or update the exercise, and any implicit assumptions will in general vary each time a new (or even the same) shock is considered. Finally, because the shocks were exogenously determined, no consideration was given to the probability of the scenarios actually occurring.

This paper briefly introduces some work done at the RBNZ towards addressing these issues. Reflecting the New Zealand banks' relatively high exposures to the housing market, the paper focuses on calculating the effect of macroeconomic scenarios on credit losses in banks' residential mortgage portfolios, but it could possibly be adapted to other portfolios (such as rural or commercial property lending). The focus on credit losses reflects the fact that they are likely to be the most significant driver of overall bank performance in stressed episodes. We

[1] The results of this exercise, along with the rest of the New Zealand FSAP, are discussed briefly in IMF (2004) and RBNZ (2004).

outline a behavioural relationship that enables us to characterise, and model systematically, the effect of a given macroeconomic scenario on each bank's credit losses, where the portfolios are characterised by the value of loans in each debt servicing ratio (DSR) and loan to value ratio (LVR) 'bucket'. This approach replaces the judgemental inputs needed from the banks in the FSAP stress testing exercise. We also introduce an approach for generating a distribution of representative macroeconomic 'states', defined by changes in house prices, unemployment and interest rates. This approach allows a simulation approach to be taken to the question of the probabilities of certain outcomes occurring. By enabling stress testers to assign probabilities to particularly outcomes, this approach represents a significant extension relative to most existing stress testing analyses, which have largely assumed away the probability aspect.[2] Section 2 reviews the existing literature on stress testing, along with some recent examples related to residential mortgage portfolios. Section 3 outlines the stress testing framework that we are working with, while Section 4 presents our model in more detail along with some initial findings. Section 5 presents the results of using our framework to replicate some of the scenarios used in the New Zealand FSAP. Section 6 concludes.

2. Existing literature

A number of international organisations, central banks, and supervisors have published papers outlining the approaches commonly used for carrying out stress tests.[3] They generally suggest that market risk stress testing is significantly more advanced than other types, followed by credit risk, with operational risk analysis very much in its infancy. But they emphasise the increasing importance of understanding these other forms of risk for financial stability purposes.

The frameworks outlined in these papers for analysing credit (and other forms of) risk all follow the same steps: 1) identify the risk; 2) construct the stress test scenario; and 3) map the transmission channels through to the banks' balance sheets. These survey pieces include a number of examples of stress tests of banks' corporate and household portfolios. However, most of these appeared to map the macroeconomic scenario to credit losses directly – missing out the middle step of modelling the effect of the macroeconomic scenarios on the borrowers' behaviour. Many of the authors cautioned against this approach because it increases the potential for unstable parameter estimates. Of those that did consider the borrowers' situations, they focused mainly on the debt servicing ability of the borrower, with little or no reference to the value of any security that lay behind the loans. Some other issues they highlighted included the fact that there are very few severe stress episodes from which to construct reliable estimates, and linear extrapolation of the relationship between macroeconomic variables and losses based on benign periods is problematic because of the suspicion that the real world is not linear – particularly in stressed situations. Finally, these papers emphasised the need to pay close attention to the cross-correlation of various risks when constructing scenarios for stress testing.

The residential mortgage component of the stress testing that was carried out as part of the New Zealand FSAP is the most comprehensive undertaken in this country to date. This 'bottom up' approach, whereby the banks individually calculated the expected losses emanating from a common set of scenarios, was also employed by the Reserve Bank of Australia (RBA) for the Australian FSAP in 2006. In addition to using a similar 'bottom up' approach as part of the Norwegian FSAP in 2005, the Norges Bank used an econometrically estimated equation to model expected losses directly as a function of macroeconomic

[2] The probability of a given outcome occurring will reflect both the number of scenarios that produce that outcome as well as the probability of those scenarios occurring.
[3] See, for example, Bank of England (2006), Bank of England (2005), BIS (2005), Sorge (2004), Jones, Hilbers and Slack (2004).

variables to perform a 'top down' stress test of the Norwegian banking system.[4] The 'top down' approach has the advantage that the exercise can be repeated easily by the supervisor for a different macroeconomic scenario. However, the Norges Bank's model was only for economy wide credit losses and could not easily be applied to different institutions with different portfolio compositions.

In 2003, the Australian Prudential Regulatory Authority (APRA) developed a 'top down' approach, whereby they applied a single model for expected losses from the mortgage book across all of the banks being analysed.[5] Each bank provided APRA with information on the structure of mortgage portfolio, including, amongst other things, the age, size, loan type, and loan-to-value ratio (LVR) distribution of their mortgage books. APRA then inputted this information into a simple equation that modelled expected losses at each bank as a function of house prices and each banks' portfolio structure. By using the same model to consider all of the institutions, it was possible to directly compare across institutions and update relatively easily. However, the assumed relationship between house prices and expected losses did not allow for explicit roles for a number of key variables, most notably the debt servicing ability of households. Hence, the model can only test for different house price changes and cannot handle more complex events, including unemployment and interest rate shocks.

In the remainder of this paper, we present a framework that is similar to the APRA approach in that it provides a systematic way of estimating expected credit losses for various scenarios across various institutions, but also extends that analysis in a number of ways. It allows for an explicit role for debt servicing and includes a macroeconomic state generating process that allows a more complex range of scenarios to be tested and is capable of assigning probabilities to different credit loss outcomes, as a (complex) function of the probabilities of the macro states.

3. The stress-testing framework

The framework we are working within is summarised in Figure 1. The starting point is gathering data on the current state of the macro economy and the financial situation of the mortgage holders. Economy wide estimates of house prices, unemployment, and interest rates are readily available from national statistical agencies. Statistics New Zealand surveys, such as the three-yearly cross-sectional Household Employment Survey (HES), the one-off Household Savings Survey (HSS) conducted in 2001, and the recently introduced annual longitudinal Survey Of Family Income and Employment (SOFIE), all provide some information on individual mortgage holders. The banks also hold data on DSR and LVR, as this information is generally provided by all borrowers at origination.

Figure 1
The Stress testing framework

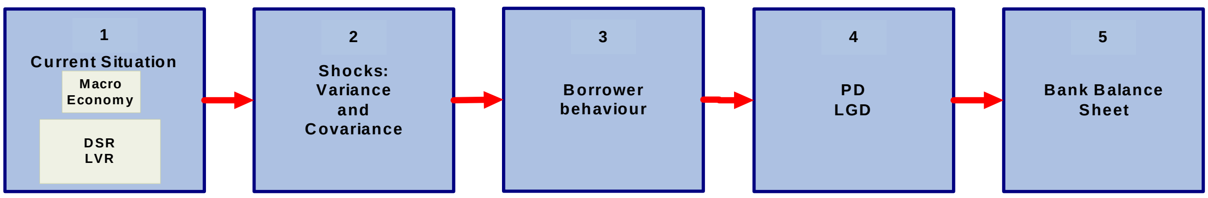

The next step is to specify the range and likelihood of possible macroeconomic scenarios in the form of a joint distribution of the three macro variables – in our case, house prices,

[4] The results of these two stress testing approaches are summarised in Hagen, Lund, Nordal and Steffensen (2005).
[5] The results of this exercise are summarised in Coleman, Esho, Sellathurai and Thavabalan (2005).

interest rates and unemployment. Carrying out this step requires identifying the means, variances, and cross-correlations of all of the relevant macroeconomic variables.

The subsequent two steps identify the effect of each macroeconomic scenario on the financial situation of the individual mortgage holders, determine the proportion of households that default in each scenario (probability of default, PD), and calculate the average size of the loss for those that do default (loss given default, LGD). In principal, the above surveys could provide information on how individual DSRs and LVRs vary with the business cycle. However, stress testing is a data intensive exercise, and in practice, even with the various surveys there are real data challenges – surveys are incomplete and infrequent and data from disparate sources are not matched – so the picture we can build up from available data is partial at best.

Having completed the first four stages, it is a relatively mechanical exercise to calculate the effect that each scenario has on the banks' credit losses. Repeating stages three through five for each of the very large number of macroeconomic scenarios identified at stage two produces a distribution of the losses emanating from the scenarios – making it possible to estimate the probability of a given loss actually occurring.

The New Zealand FSAP stress testing exercise can be considered as single draws of the macro state within this framework, coupled with bank judgement for the behaviour of borrowers. Again, the fact that they considered only a limited number of scenarios precluded them from ascertaining the probability of the estimated losses actually occurring.

4. The framework in more detail

The behavioural relationship we have developed provides a systematic approach for determining the effect of the macroeconomic state on household finances and behaviour, and thus banks' expected credit losses. The simple intuition behind the relationship is that a borrower defaults when two conditions are met simultaneously. The first condition is that the borrower can no longer afford to service the loan, and is forced to sell their house (a stressed sale). The second condition is that the value of the house is less than the value of the loan. If both conditions are met, the bank takes ownership of the security – the house. The bank then sells the house and incurs a loss equivalent to the shortfall between the loan and house value and additional transaction costs relating to the disposal process.

If only one of these conditions holds, there is no default event. If the borrower can afford to service the loan, then they will continue to do so even if the value of the house falls below the value of the loan. If, on the other hand, the borrower cannot afford to service the loan and the house is worth more than the loan value, then the borrower sells the house and repays the loan in full. Again there is no loss to the bank. The key assumption here is that borrowers do not default 'strategically'. That is, they do not base a decision to default solely on whether the value of the loan is more than the value of the house. The assumption that borrowers will continue to service a loan when they have negative equity is critical to the model structure and contrasts with much of the US literature on the determinants of housing defaults. It is important then to set out the reasons why this assumption was adopted.

First, it captures a lot of actual behaviour in the jurisdictions where there have been major falls in housing prices. In the UK and Hong Kong, for example, borrowers tended to hang on to houses if they could afford to service the debt even when they had substantial negative equity.

Second, we think that this is how New Zealanders would behave in the face of a large fall in house prices that put them in a negative equity position: New Zealanders' strong commitments to their family home; a psychological unwillingness to acknowledge a loss by crystallising it in a sale; and the fact that debts are not discharged if the lender forecloses on the security, are factors supporting this contention. While rental housing investors could behave differently, the core model is designed to deal with the conventional owner-occupier loan portfolios where the above factors will be at their strongest. The possibility that investors could behave differently can be addressed in a purpose built module for investment loans.

Third, while there is likely to be some feedback effect from a negative equity position to the borrower's willingness (as opposed to capacity) to service the loan, this effect is likely to be of a second order of importance. As we had no way to ascertain the size of any such effect we thought it was better to leave it out rather than put in place a formal structure that accommodated the negative equity effect that had little or no empirical content.

The model is structured to calculate the distribution of loss outcomes given this joint default requirement. It first calculates the proportion of households that will experience a stressed sale event for a given combination of interest rates and unemployment outcomes.[6] A stressed sale can be due to an idiosyncratic event (ie specific to a borrower) such as a marriage separation or a case of financial mismanagement, or to systemic economic events (affecting all borrowers), such as increases in interest rates and decreases in income. In this version of the model, we are only modelling wage and salary earners, and therefore proxy income using the unemployment rate. For a particular change in interest and unemployment rates, the model uses a behavioural relationship to calculate what proportion of the portfolio will be subject to a stressed sale. We explain this relationship in more detail later in this section.

The next step is to determine whether there will be a default on the loan as a result of the house being sold. The model calculates the value of the house security for each loan in the portfolio drawing from a conditional distribution of individual house price changes. This distribution captures the extent to which the price of the borrower's house has changed by more or less than the movement in aggregate house prices given by the macro state generator. A comparison between these individual prices and the loan value determines whether the house is sold by the borrower, or whether the bank forecloses.

In the event of foreclosure, several other factors, in addition to the initial gap between the value of the house in the hands of the owner and the loan, determine the extent of the economic loss to the bank. They include disposal costs, a 'foreclosure discount' (a further reduction in the value of the house because the bank has foreclosed in a downturn environment) and an allowance for the fact that it takes time to sell the house and receive the proceeds. Once the individual losses are calculated, the losses are then added together to obtain total portfolio loss for the scenario.

This process is repeated using 10,000 random draws for the macro states, generating the entire credit loss distribution. The model can be used to calculate loss distributions for individual banks, where the only information required from the banks are the values of loans in each LVR/DSR bucket. This enables direct comparison of the relative riskiness of different banks' loan portfolios. It is also possible to use survey results for the economy wide DSR/LVR structure to generate default and loss rates for the entire banking system.

---

[6] Although this discussion is in terms of the entire loan portfolio, the same principal can be applied to subsets of the portfolio, where those subsets are characterised by the DSR and LVR of the mortgages.

Within this framework, there are two key areas that require further discussion: the macroeconomic model for generating the outcomes for house price inflation, the change in interest rates and the change in unemployment; and identifying the behavioural relationship between the macroeconomic variables and the proportion of the portfolio that will be subject to a stressed sale.

Modelling the macroeconomic variables

If the framework is only being used to determine the effect of a particular scenario, then it is not necessary to have a systematic macro state generator to produce that scenario – a single scenario is exogenously determined and put into the framework to calculate the expected credit losses emanating from that scenario. However, a macro state generator is required to calculate the loss distribution, and hence the probabilities.

One option is to use a traditional econometric technique, such as a vector autoregression (VAR), estimated over a long period of historical data. While the VAR methodology is potentially useful in some areas of economics and finance, there are some factors that make it less useful for our purposes:

- Data sets that are long enough to capture tail events – which are the most interesting for stress testing purposes – are generally susceptible to parameter instability resulting from structural change.
- VARs are generally linear, and focus on the average relationship between the relevant variables. There is a risk that they may understate, or miss altogether, economic relationships that only become apparent in stress events.

For these reasons, we are developing a calibrated macro state generator that relates house prices, interest rates and unemployment. The generator is static in nature, with each scenario representing a single three year period. We see three years as long enough to capture most of the effect of the various scenarios. The macro-variable shocks are drawn from a multi-variable normal distribution. The assumption that the macro-variables are normally distributed does not necessarily represent a view that this is the best depiction of reality. Rather it is adopted because the model becomes more analytically tractable and because it is the conventional risk language metric. It makes it easier to explain levels of risk and changes in that risk.

The parameters of this joint distribution (ie means, standard deviations, and cross-correlations) are necessary and important inputs for determining the probability of various outcomes. The ‘true’ values of these parameters are unobservable, and obtaining appropriate estimates for them requires considering a number of different avenues. This is one of the areas of the framework that we are still developing, and one of the benefits of the approach we have taken is that alternative parameter estimates can easily be used. We now spend a few moments discussing the approaches that we have used in arriving at our current parameter values.

The mean change in interest rates and the unemployment rate has been set to zero, reflecting the assumption that these variables are mean reverting. House prices are assumed to grow at around or slightly above medium term CPI inflation.

In terms of the standard deviations, we considered historical New Zealand data and the number and severity of house price crashes observed internationally. In a survey of house price booms and busts for 14 OECD countries over 30 years, Helbling (2004) found that ‘busts’ were relatively frequent (20 events with a real house price fall of more than 14 per cent) and large (the average real fall in prices was 27 per cent). Empirically estimated standard deviations for New Zealand house prices implied lower frequencies for such events. As we saw no reason to expect New Zealand to have fewer house price crashes than these other countries (small countries tend to be at least as volatile as large countries), we calibrated the house price volatility to be about 30 per cent higher than the historical New Zealand data estimate. The assumed interest rate volatility of 2.5 percentage points reflects a margin over the international volatility figure implicit in the interest rate shock data collated by Helbling. The margin reflects our judgement that, as a small heavily indebted country, New Zealand will tend to exhibit more interest rate volatility than the OECD average. The unemployment rate volatility reflects a judgement that a 6 percentage point increase is a plausible outcome for a two standard deviation stress event.

Before turning to the cross-correlations, it is useful to reiterate that the events that we are most interested in are in the tails, because that is where the largest credit losses occur. Therefore, we have chosen values for the cross-correlations that we feel most accurately account for the likelihood of stressed events. As a result, our calibration is likely to understate the frequency of very small losses, but hopefully be relatively good at capturing the frequency of large losses (Figure 2).

Figure 2
Stylised loss distributions

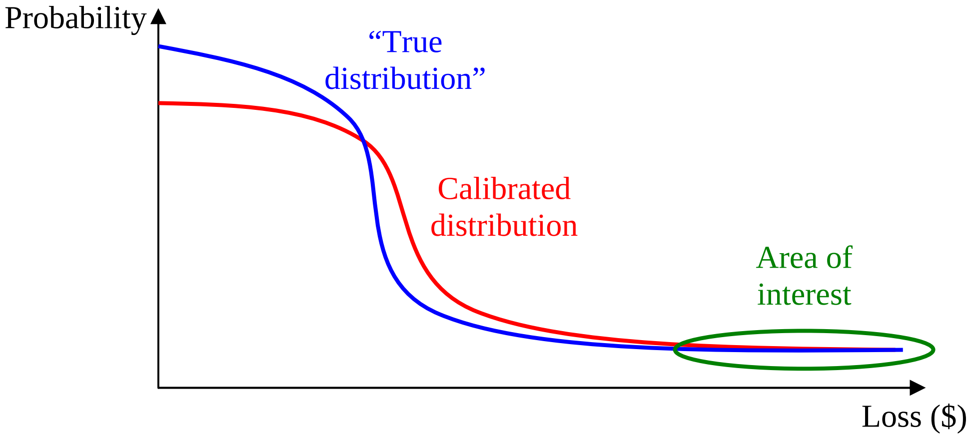

Increases in house prices are typically associated with lower unemployment, and a decrease in unemployment is typically associated with higher house prices, all other things equal. This suggests that the cross-correlation between these variables is likely to be quite negative. This is generally confirmed by historical data, but the exact value fluctuates quite a lot depending on the sample period.[7] We have currently opted for a value of -0.5 for the cross-correlation during stressful events.

The interest rate / house price cross-correlation is more complicated because the direction and magnitude of the correlation is likely to be dependent on the source of the shock. For example, a positive domestic demand shock is likely to see house prices and interest rates rise. Conversely, a cost shock is more likely to put upward pressure on interest rates and downward pressure on house prices. This is demonstrated by the fact that the estimated cross-correlation for New Zealand fluctuates either side of zero depending on the sample period. Using the average value in the macro generator would most likely generate to few high interest rate / low house price and low interest rate / high house price outcomes, and therefore

[7] Real values were used because we felt high inflation prior to 1992 would mean that estimated cross-correlations between nominal variables would overstate the likely cross-correlations during more moderate inflation periods.

produce too few very high and very low credit loss outcomes. We have chosen to use a negative cross-correlation between house prices and interest rates because we feel this is the most appropriate calibration to generate the fat tail of the joint distribution of macro states that we are interested in (ie the ones that produce large credit losses). In order to determine how negative to set the cross-correlation, we considered some IMF work that showed that house price crashes were typically immediately preceded by interest rate increases (Helbling 2004). However, it is possible to point to international episodes in history where large house price declines have not coincided with higher interest rates in that country, which suggests that the cross-correlation is probably not as negative as the cross-correlation between house prices and unemployment.

It is likely that the cross-correlation in a small open economy such as New Zealand would be more negative than it is in the US. In a large country like the US, monetary policy there will typically only tighten in the face of falling house prices if there is a genuine risk of inflation increasing. Conversely, to the extent that international interest rates have a significant influence on longer-term (1 to 5 year) interest rates in New Zealand – an area of the curve where many New Zealand banks fund their mortgage book – then it is quite conceivable for the New Zealand central bank to lower short-term interest rates in response to a domestic slow down and house price slump, but longer-term interest rates to remain high or even increase in response to high global interest rates: the inverse of the situation seen in New Zealand through 2005. A domestic confidence shock hitting New Zealand, which was large enough to increase the premium on interest rates in New Zealand relative to the rest of the world, is another example of a potentially negative correlation between house prices and interest rates in New Zealand that a large economy like the US would be significantly less susceptible to. Similarly, a global 'flight to quality' could quite possibly see New Zealand interest rates rise and US interest rates fall even if (and potentially related to) the New Zealand economy entered a slowdown. The current calibrated values are shown in Table 1.

Table 1
Calibrated macroeconomic parameters

	Mean	Standard deviation	Cross-correlation	
			Unemployment	Interest rates
Unemployment[a]	0	3.0	–	–
Interest rates[b]	0	2.5	+0.3	–
House prices[c]	7.5	17.5	-0.5	-0.3

Notes: a) Cumulative percentage point change in unemployment rate over three years
b) Cumulative percentage point change in interest rates over three years
c) Cumulative per cent change in house prices over three years

The behavioural core of the framework – determining the probability of a stressed sale

In this section, we outline the model that links the macroeconomic scenarios to banks' losses in a systematic way. The core of this part of the model is the behavioural relationship between the macroeconomic variables and the probability a household experiences a stressed sale event.[8] This is the same as modelling the proportion of households that experience a stressed sale event in response to a given macroeconomic scenario. As outlined above, a stressed sale can be due to an idiosyncratic event such as a marriage separation or a case of financial mismanagement, or to a systematic economic event that increases interest rates and / or unemployment (Equation 1). The probability of a household experiencing an idiosyncratic or systematic event severe enough to force them to sell their house will be an increasing function of that household's initial DSR.

[8] Note that here we are only calculating the probability of stressed sale. As noted at the beginning of Section 4, this information is combined with the idiosyncratic house prices to calculate actual losses.

$$\Pr(SS_{sj}) = f(\beta, \Delta i_s, \Delta u_s, DSR_j) \qquad (1)$$

where: PR(SS_{sj}) is the probability a household j experiencing a stressed sale given scenario s; β is the proportion of households that are expected to experience a stressed sale event due to idiosyncratic factors; i_s, and u_s are respectively the economy wide interest and unemployment rates that characterise the macroeconomic scenario; DSR_j is the DSR for household j prior to the shock occurring.

While ideally we would have a rich enough data set to estimate this function directly, this dataset does not exist. The period over which banks have good mortgage default data is very short and only covers a period of relatively benign macroeconomic outcomes. For this reason, we have calibrated our model using a combination of current, historical and international experience, international research, and expert judgement.

Our estimate of the idiosyncratic effect was drawn from a range of recent information provided by New Zealand banks. In calibrating the impact of the systematic factors, we have relied heavily on the British housing crisis of the early 1990s. This required a large number of steps. Our starting point was reported Building Society loss data relating to that period. This was complimented with insights from UK research that encompassed this period (in particular, Whitely *et al* 2004) and some judgements about the distribution of borrowers' debt servicing ratios gained from our analysis of similar New Zealand portfolios. Using all of this information, we were able to obtain an approximation of the behavioural relationships between changes in interest rates and unemployment.

As is illustrated below, we were able to confirm that these behavioural relationships are relevant to New Zealand by using the model to repeat the relevant FSAP scenarios, and compare the results to those produced by the banks in 2003.

Some interesting insights

Combining the macroeconomic state generator and the behavioural relationship discussed above enables us to draw some interesting insights about individual scenarios and the loss distribution. A sample of some of the more important observations implied by the model are as follows:

- The biggest driver of risk is a simultaneous interest rate increase and a house price fall. The volatility of interest rates and house prices and the way they are correlated are, therefore, the biggest determinants of expected losses in stressed environments. Changes in unemployment are less important. The simple intuition here is that unemployment only affects borrower's who are unemployed, whereas every borrower is affected by an increase in interest rates.

- The duration structure of a mortgage portfolio can be an important determinant of risk. Floating rate mortgage portfolios are riskier than their fixed rate counterparts, and the longer the average fixed rate term the better.

- The loss distribution is clustered close to the zero mark but has a long tail. Thus, it is possible for a risky portfolio to experience considerable periods of very low realised losses, while still remaining susceptible to sizeable losses in the event of a macroeconomic shock.

5. FSAP stress test replication

In this section, we use our model to replicate the two New Zealand FSAP stress tests related to residential mortgages. The first was a single factor test. Banks were required to calculate losses in their housing portfolios due to a stress event that combined a 20 per cent drop in house prices and an increase in the unemployment rate from 5 to 9 per cent. In this scenario, banks reported aggregate credit losses equivalent to about 1.1 per cent of their residential mortgage assets.

The second scenario was a more complex event that modelled a shock to New Zealand's external credit rating. It involved a substantial shock to short term interest rates, a fall in the exchange rate, a decline and then partial recovery of house prices and an increase in unemployment. Credit losses for the scenario were about 1.0 per cent of total housing loans.

The DSR and LVR structure of the New Zealand banking system residential loan portfolio assumed for this exercise came from the 2004 HES.

Using the model discussed in this paper, a simulation of the first shock generated an average portfolio loss of 0.6 per cent and the second a loss of 1.2 per cent. Given an array of technical differences between the FSAP exercise and our simulation, the loss rates are reasonably close. The higher results for the first scenario reflect banks' assessments that unemployment will have a greater impact than our model suggests.

The second scenario mainly captures the effect of the interest rate increase on borrowers' servicing costs. Here the closeness of the results suggests similar assessments of the underlying borrower behaviour. This is a useful independent confirmation of the reasonableness of the key interest rate /stressed sale relationship in our model.

A second test was to run the second scenario again but with an increase in the assumed house price fall – a fall of 15 per cent when house price had just increased by similar magnitude is arguably not much of a 'stress' test. Increasing the assumed decline to 25 per cent increases the loss rate from 1.2 per cent to 2.1 percent.

Finally, the loss distribution generated by the model was used to put likelihoods on the sizes of the losses generated by the stress scenarios. Table 2 expresses the size of each shock in terms of the probability that a loss of that size, or greater, will occur given the model's macro-economic and behavioural structure. The first stress test could be described as a 'one in twelve event' while the second is a 'one in twenty'.

Table 2:
Implied probability of equal or larger credit losses

	Likelihood (per cent)
Scenario 1	8.2
Borrowing shock:	
15 per cent house price fall	4.6
25 per cent house price fall	2.0

Section 6. Conclusion

Stress tests are one of the important tools used by supervisors and central banks for assessing the stability of financial systems. This paper has presented an approach for systematically modelling the affect of changes in house prices, interest rates and unemployment on banks' credit losses. This approach requires very little input from the banks, and therefore enables the stress tester to quickly and easily consider alternative macroeconomic scenarios in a consistent manner across a large number of institutions. In addition, this paper has presented a method for generating a distribution of representative macroeconomic 'states', defined by changes in house prices, unemployment and interest rates. We think this represents a considerable advance on the existing stress testing literature in that it enables the stress tester to apply probabilities to various loss outcomes actually occurring.

The framework presented today lays out the foundation for further work. For example, we are still experimenting with the specification and parameterisation of the model that relates the macroeconomic variables to the probability of stressed sale. In addition, we are investigating other avenues for determining the most appropriate values for the variances and the cross-correlations in the macro-state generator. Using the appropriate values for these parameters is crucial for determining the loss distribution. The main sources of information for these tasks are historical housing market collapses, but we are also interested in other possible guides. We are also testing the sensitivity of the results to different parameter choices.

Once we are comfortable with the core model, we hope to carry out and present some further stress testing of residential mortgage portfolios. The current model can also be extended to allow for fixed rate mortgages, investor housing and self employed. In addition, it is an option to try to extend the model to consider other portfolios, including agricultural and commercial property. Finally, it would potentially be very interesting, but difficult, to develop a dynamic version of the model that allows the analyst to observe the various scenarios unfold period by period.

References

Bank for International Settlements (2005), "Stress testing at major financial institutions: survey results and practice," Report by a working group established by the Committee on the Global Financial System.

Bank of England (2006), "Systematic stress testing," Box 6, *Bank of England Financial Stability Review*, July, pp 45–50.

Bank of England (2005), "Stress testing as a tool for assessing systemic risks," *Bank of England Financial Stability Review*, June, pp 116–126.

Coleman, A, N Esho, I Sellathurai and N Thavabalan (2005), "Stress testing housing loan portfolios: a regulatory case study," Paper presented at the *Basel Committee on Banking Supervision Conference* on "Banking and financial stability – a workshop on applied banking research," Vienna, April 20–21.

Hagen, J, A Lund, K Bjorn Nordal and E Steffensen (2005), "The IMF's stress testing of the Norwegian financial sector," *Norges Bank Economic Bulletin*, 4.

Helbling, H (2004), "Housing price bubbles – a tale based on house price booms and busts."

International Monetary Fund (2004), "New Zealand: financial system stability assessment," *IMF Country Report*, No 04/126.

Jones, M, P Hilbers and G Slack (2004), "Stress testing financial systems: what to do when the Governor calls," *IMF Working paper*, WP/04/127.

Reserve Bank of New Zealand (2004), "Stress testing the New Zealand banking system," *Reserve Bank of New Zealand Financial Stability Report*, October, pp 27–31.

Sorge, M (2004), "Stress testing financial systems: an overview of current methodologies," *BIS Working Papers*, No 165.

Whitely, J, R Windram and P Cox (2004), "An empirical model of household arrears", *Bank of England Working Paper*, No 214.