Reserve Bank of New Zealand

Project TUI

A structural approach to the understanding and measurement of residential mortgage lending risk

March 2008

Ian Harrison
Chris Mathew

Abstract:
TUI is a structural model of the residential housing loan default process. It is designed to investigate major loan loss events, and to calculate economic capital for housing loan portfolios, in jurisdictions where empirical evidence of tail event losses is slight or non-existent. It works by first describing the key risk characteristics of a small set of macroeconomic variables which are relevant to the housing loan default process. Macroeconomic scenarios are drawn from a joint distribution of these variables and bank and borrower behavioural modules are used to generate a loan loss rate for each scenario. A large set of loss outcomes generates a loan loss distribution which is used to calculate, amongst other things, capital requirements, average default rates and stressed loss given default rates by risk bucket and for overall portfolios. The model has been used by the Reserve Bank to assess the applicability of the Basel II IRB housing capital model to New Zealand conditions and to provide a benchmark for assessing New Zealand IRB banks' housing capital models.

Contact : Ian Harrison is a Special Adviser in the Prudential Supervision Department of the Reserve Bank of New Zealand . Contact details are e-mail harrisoni@rbnz.govt.nz or ph. 64-4-471-3756.

Ref #2747610

Contents

1 Introduction and overview

TUI is a structural model of the residential housing loan default process. It is designed to investigate major loss events in residential housing loan portfolios that may have occurred rarely, if at all, in many jurisdictions. These events are nevertheless possible and are relevant to questions such as the amount of capital a mortgage lender should hold or what might happen in a particularly acute stress event. The name of the model – **T**ool for **U**nobserved-event **I**nvestigation – TUI, captures this focus on the analysis of 'tail-end' events when there is little, if any, reliable data from actual events.

Obvious uses for this kind of model are in the Basel II IRB pillar 1 validation and pillar 2 assessment processes which potentially pose some major challenges for supervisors.

For example, in the New Zealand context, there is a question as to whether the IRB housing capital model has been appropriately calibrated to measure risk in New Zealand housing loan portfolios. Does one size[1] fits all or is there something about New Zealand's circumstances that makes housing loan portfolios more or less risky than the international norm?

Even if the underlying IRB model is robust its implementation raises some difficult data problems. In the absence of good long run data series and often a complete lack of data on the kind of events that generates tail losses, how does the supervisor know whether the key inputs into the IRB model – the probability of default (PD) and the stressed loss given default (SLGD) – are near the mark? In the absence of long run historical data, banks have tended to rely on relatively short data sets implicitly assumed that, with some moderate upward adjustment, these provide an unbiased estimate of the true default and loss rate averages. Does this kind of approach generate reasonable estimates?

TUI provides an independent reference point to help address these kinds of questions. It combines an explicit structure of the loan default and loss process with estimates of behavioural and macroeconomic risk driver coefficients to produce a distribution of loss outcomes. This is used to generate both IRB model input estimates and risk weights, allowing a comparison between a fuller structural default model and the more parsimonious Basel II model outputs.

TUI can also produce answers about mortgage portfolio losses in particular stress events. Given a set of assumptions about the mortgage interest rate, unemployment and house price changes it will generate internally consistent loss outcomes which make it easy to compare one scenario with another. At present the model has been optimised to deal with the far tail events rather than more moderate stress events. However, it handles a reasonable range of scenarios and the structure can be further developed to handle a wider range of events and produce more precise loss estimates.

While the full model structure appears to be complex, the core of the model is designed to be transparent and accessible. It does not rely on arcane theory or intricate mathematical formulations. It should be able to be understood and used by anyone who can deal with a few simple algebraic formulas.

[1] For example the IRB housing equation has a common correlation coefficient of 0.15 for all housing exposures.

The simple intuition behind the structure of the model is that a mortgage loan default occurs when two conditions are met simultaneously. The first is that the borrower can no longer afford to service the loan. The second is that the value of the house is less than the value of the loan. If both conditions are met the borrower 'puts' the security – the house – back to the bank. The bank then sells the house and incurs a loss equivalent to the shortfall between the loan and net house sale proceeds.

If only one of these conditions holds there is no default event. If the mortgagee can afford to service the loan then he will continue to do so even if the value of the house falls below the value of the loan. That is, the borrower does not act to increase his net wealth by putting the asset back to the bank. If, on the other hand, the borrower cannot afford to service the loan, but 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 assumption that borrowers will continue to service a loan when they have negative equity is critical to the TUI 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 the TUI assumptions were adopted.

First, it captures a lot of actual behaviour in the jurisdictions where there have been major falls in housing prices. In the United Kingdom 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 commitment 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 TUI 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 better to leave it out rather than put in place a purely 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 the joint default requirement. The model first calculates the probability that there will be a stressed sale event. A stressed sale occurs when the owner has to either sell the house or put it to the bank because she can no longer service the loan. This can be due to an idiosyncratic event such as a marriage separation or a case of financial mismanagement, or to a systematic economic event which increases interest rates (and hence debt servicing costs), and/or unemployment.

The model runs a Monte Carlo process which generates a range of macroeconomic scenarios each of which is described by specific values for the floating mortgage interest rate, unemployment and the house price index. The probability that a particular scenario will be drawn is described by a joint distribution of the changes in the interest rate, the house price index and unemployment rate. This joint distribution is generated by volatility and correlation inputs into the model.

For each macroeconomic scenario the model then uses a behavioural relationship to calculate the proportion of the portfolio which will be subject to a stressed sale. The next step is to determine whether there will be a default. Taking the house price index value for the scenario as the starting point the model determines the value of the house security for each loan in the portfolio (by drawing from a distribution of individual house prices). A comparison between the individual house prices and the loan value determines whether the house is sold by the borrower, at no loss to the bank, or whether the bank forecloses and the event is recorded as a default. For each case of default the model then calculates the loss given default given the gap between the loan and the net realisation from the house sale and associated costs. These losses from the 10,000 Monte Carlo simulations will describe the loss distribution.

There are, then, four main modules in the basic model. The first is the macro-economic module which generates the model's systematic risk. It describes the volatility and correlation structure of the variables which impact on borrowers' servicing capacity and on the general level of house prices. The second describes how demographic factors impact on the stressed sale rate. The third component is the behavioural module. It describes how a portfolio of borrowers reacts when their capacity to service their loans is affected by changes in loan servicing costs and unemployment. The fourth is the loan loss module which determines whether there is a default and the extent of the bank's losses when there is a loan foreclosure.

All of the model coefficients can be readily changed to reflect different judgements about the structure of the economy or of borrowers' behavioural characteristics.

Once a loss distribution has been calculated the model can generate an array of outputs including the long run probability of default, average and downturn loss given defaults and risk weights for an overall loan portfolio and for individual risk buckets (defined with respect to the loan to valuation ratio (LVR) and the borrower's debt servicing ratio (DSR)).

TUI's PD and LGD estimates cannot be compared directly with IRB consistent model estimates because the models use different definitions of default. TUI uses an economic definition of default (where each default is a solvency event and results in a loss) while the New Zealand IRB models use an accounting definition. As any loan which is 90 days past due is counted as being in default, the IRB definition picks up a lot of additional liquidity events.

A second model, called Basel-TUI, is required to translate the TUI outputs into IRB consistent outputs.

Figure 1
Model Structure

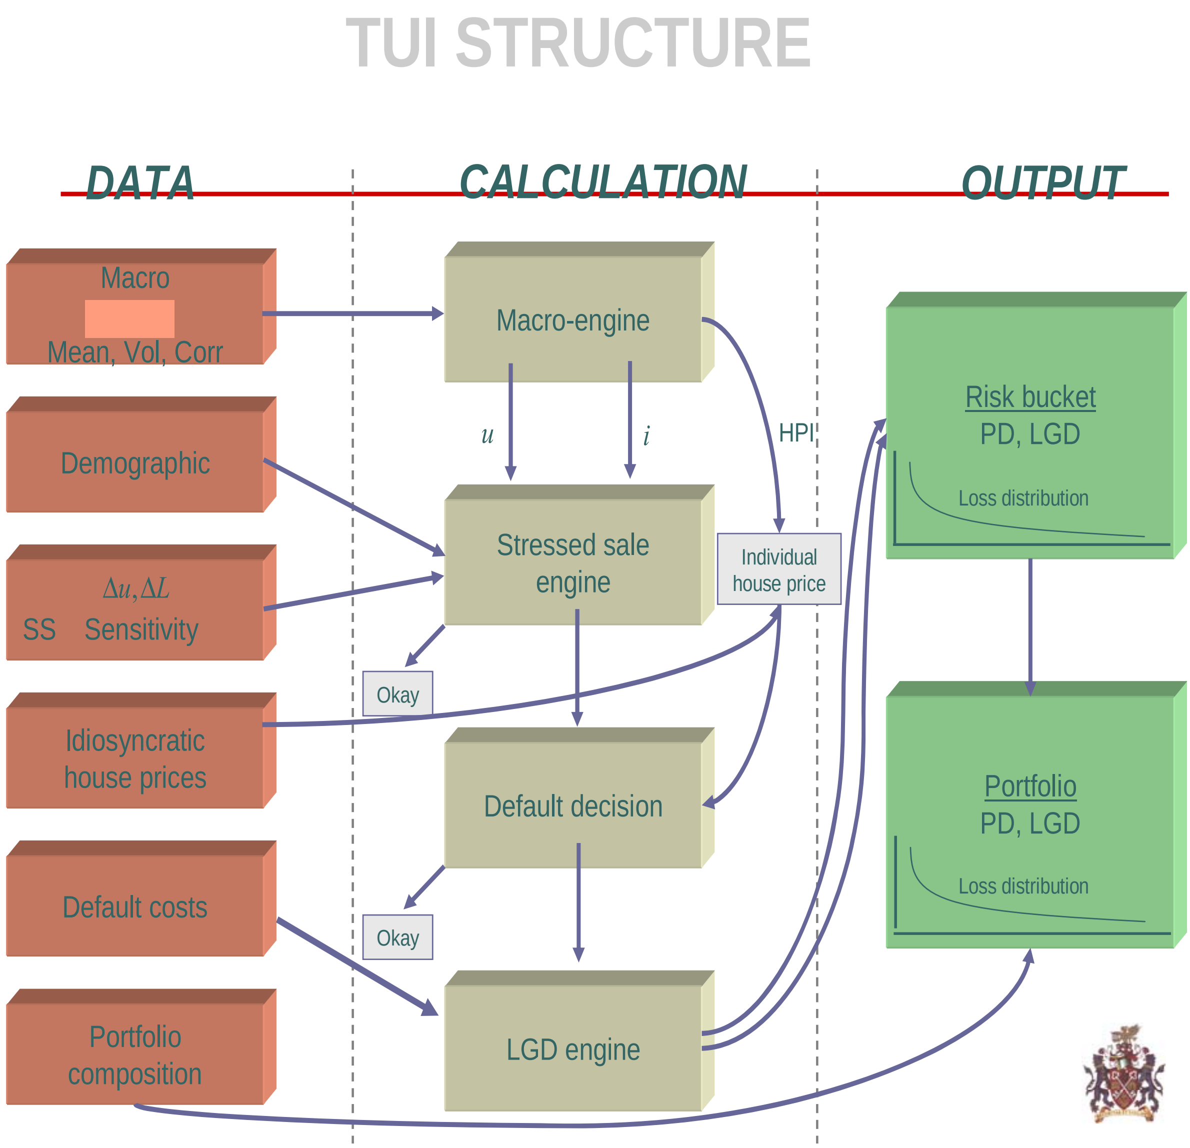

The base model, which is used to investigate and describe the main properties of housing lending risk, assumes that all borrowers are wage and salary earners rather than self employed and that all borrowing is at a floating rate. However, TUI also has a number of secondary modules which can be used to address a wider range of loan characteristics and borrower circumstances. Questions such as the effect of mortgage insurance, the impact of fixed rate as opposed to floating rate mortgages, and of self employment rather than salary and wage status, can be investigated.

As well as investigating the properties of the Basel II IRB model TUI can also be used in stress testing mode.

The model's quantitative outputs are generally subject to a reasonable margin of variation and undue credence should not be placed on the precise quantitative outputs. Many of the outputs are, nevertheless, qualitatively robust. They have the right signs, the right orders of magnitude and are economically plausible.

The following is a sample of some of the more important, and not always obvious, TUI results.

- 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 capital requirements.
- The duration structure of a mortgage portfolio is 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 lower is the risk.
- Residential mortgage lending appears to be substantially more risky than Basel II modelling by banks would suggest and some higher risk risk-buckets may require more capital than required by the standardised model.
- Low observed default rates in benign times can be consistent with a risky portfolio and a high capital requirement.
- TUI results are consistent with the stylised facts of several housing stress events that have recently occurred around the world.
- The Basel II IRB housing lending model may require recalibration for New Zealand conditions. This largely reflects the greater exposure of floating or short term fixed rate mortgage portfolios to systematic risk.

The paper is organised as follows:

Section two sets out the structure of the core model which applies to employed, owner-occupier borrowers with floating rate loans.

Section three describes how the macroeconomic, behavioural and loan loss inputs into the model were derived.

Section four presents the core model outputs. These include PDs, average and stressed LDGs and risk weights.

Section five presents a sensitivity analysis of key coefficients and other model inputs and an assessment of the stability and robustness of the model's outputs.

Section six compares the TUI outputs with the US Federal Reserve's and APRA's modelling results.

Section seven describes the add-on modules which can be used with the core model. They cover: fixed rate mortgages; mortgage insurance; loans to the self employed and 'rational' forbearance.

Section eight applies TUI to Basel II issues. The applicability of the Basel II IRB housing model correlation coefficient assumption to New Zealand conditions is discussed and the structure of the model – Basel-TUI - which converts the TUI PD and LGD outputs to an IRB consistent form is described.

Section nine presents the results of TUI simulations of a number of historic and simulated housing loss episodes. They include New Zealand loss experiences in current good times and during the interest rate shock of the 1980s; the bank generated housing losses for the 2004 New Zealand FSAP exercise; Australian loss experiences of the 1980s and early 1990s; and the Hong Kong experience of the late 1990s.

The final section provides an assessment of the usefulness of the model.

2 The Core TUI Model

The first part of this section discusses three key structural features of the model: the assumption of normality in the distribution of the macroeconomic variables; the model's time horizon, and the definition of default. The second and third part sets out the structure in more detail and the formal model specifications.

2.1 Three key structural features

Distribution of macro-variables

It is assumed that the macroeconomic variables are normally distributed. This does not necessarily represent a view that these variables are normally distributed but was adopted for the following reasons. First, it makes for a more analytically tractable model. Second, normality is the conventional risk language metric. It makes it easier to explain levels of risk and changes in risk by working with this convention. Third, it preserves comparability with the Basel II IRB model which assumes that the value of underlying asset, which backs loan portfolios, is normally distributed. If we used a different distribution with a fatter tail this would be equivalent to adopting a tougher solvency test than the IRB model assumes.

If the user were to prefer an alternative distribution assumption then this can be done in an ad hoc way by adjusting the macro variable volatility assumptions.

Definition of Default

The definition of default in TUI is not the same as the Basel II IRB definition. For practical purposes the Basel II IRB definition of default is whether a loan is 90 days or more past due. This is an accounting definition which may not necessarily be consistent with the theoretical structure of the IRB housing model. The underlying model assumes that every default event generates a loss, or, in other words, is a solvency event. The 90 day past due definition, however, also captures liquidity events which do not generate material losses and which are not, therefore, relevant to the calculation of a bank's risk and capital requirements.

In recent years New Zealand banks have been reporting around twenty liquidity events for every solvency event in their default data. The introduction of a dominating mass of economically irrelevant liquidity events into the data set for estimating PDs and stressed LGDs has two effects. First, it alters the size of the average PDs and LGDs and changes their distributions. This makes it more difficult to compare the numbers with data based on more conventional, economic, definitions of default. An economic LGD, for example might be

30 percent whereas the equivalent IRB measure might be only 3 percent. Second, it tends to mask the relationship between the key LVR risk driver and regulatory risk weights.

To avoid these problems default, in the TUI framework, is defined in economic terms. Every default generates an economic loss.

Time horizon

The solvency standard of 0.999 used in the Basel II framework relates to a one year horizon. TUI uses a three year horizon. The decision to use a three year horizon was based on two considerations. The first is the length of housing shocks. It typically takes a number of years before an economic shock is fully reflected in housing lending loss outcomes and it doesn't make much sense to attempt to decompose longer horizon events, into sets of one year of events, to fit into a one year horizon model.

The second is a view on the policy purposes of capital. From this perspective capital can be viewed as a buffer to absorb losses which occur while a shock is being recognised; plans formulated, and new capital raised. A three year horizon would reflect the view that the appropriate test is that bank capital should be adequate to cover three years of losses with a defined probability.

Note that the three year horizon should not be taken literally. Rather it is a stylised convention which might, in some circumstances, compress and conflate calendar events.

Note also that while the Basel II IRB model uses a one year horizon this need not in itself represent a view on the appropriate horizon for the analysis of credit events. It is simply a measurement convention. The underlying theoretical model does not have a time dimension to it and can work over any time horizon.

To preserve comparability with the Basel II IRB outputs the solvency standard has to be adjusted for the three year horizon. This was done by first generating a capital requirement using TUI with the IRB solvency standard and with one year PDs and LGDs. The PD is then multiplied by just under three to obtain a three year PD. The capital amount generated for the one year horizon, the three year PD and the LGD are entered into the IRB housing equation which was solved for the solvency standard. While this figure will vary according to the original one year PD it was found to be reasonably stable over a range of the most relevant PDs. A common solvency standard of 0.987 has been used for all risk buckets.

2.2 Model structure

The intuition behind the structure of the model is that a residential mortgage loan default occurs when two conditions are met simultaneously. First, the borrower can no longer afford to service the loan. Second, the value of the house is less than the value of the loan.

To determine whether there is a default and the probability of this occurring the model works through a two stage process. The first step is to calculate the probability of a stressed sale (PrSS). A stressed sale is defined as a sale which occurs when the owner is forced to sell the house or put it to the bank. The stressed sale rate is determined by two sets of drivers. The first are demographic which includes idiosyncratic events such as divorce, sickness and financial mismanagement which are captured by the demographic coefficient. The second set of drivers are systematic. Changes in the unemployment rate and the mortgage interest rate will impact on borrowers' capacities to service their mortgages.

The second step is to determine, for any stressed sale event, whether the value of the loan is greater than the net value of the house for each loan in a portfolio. If the loan value is higher the loan will go into default. Default will be determined by the change in the house price index, the change in the value of a particular borrower's house, and by the original LVR of the loan.

Once the proportion of loans in default in a portfolio is calculated, the model then calculates the loss given default given the house price outcome and a number of transaction cost variables. This process is repeated 10,000 times with different macroeconomic draws to generate a loan loss distribution.

Choice of main risk drivers

The model focuses on two key risk drivers. They relate leverage of the loan – the current loan to valuation ratio (LVR) – and the capacity of the borrower to service his loan, which is captured by his current debt servicing ratio (DSR).

A mortgage portfolio is divided into a set of risk buckets which are defined by their LVRs (loan to valuation ratio) and DSRs (debt servicing ratio). All of the model processes are applied at the risk bucket level. Overall portfolio numbers are generated at the end of the process as a value weighted average of the individual risk buckets results.

There are a number of other factors that may also drive losses in residential mortgage books but these are not modelled explicitly. Some of these factors are:

- The elapsed time from the origination of the loan. Several studies have found a stable relation between time from origination and the probability of default. Default rates are generally low to start with; peak after two or three years and then tail off.
- The borrower's absolute income level of income.
- The 'quality' of the borrower. In the US this is captured by formal models which score past behaviours which are related to the probability of default.
- The size of the loan. Some studies have found higher losses on large loans because they are secured against 'luxury' dwellings whose prices suffer disproportionately in a downturn. There is also evidence of higher loss rates on small loans at the bottom end of the market.

These factors were not modelled explicitly because:

- They would make the model extremely complex. A matrix of 6 x 6 DSR and LVR risk categories has 36 risk buckets. Adding just four income size categories and five credit categories, for example, would generate a total of 720 risk buckets.
- The relevant data is either not available or is of indifferent quality.
- The theoretical impact of the risk driver on losses may be poorly understood and difficult to model.
- It is not necessary to model these risk drivers explicitly if it can be assumed that they are evenly spread between the risk buckets defined by LVR and DSR. The impact of these variables will then be subsumed within the average behavioural coefficients of the individual risk buckets.

Some of these considerations do not appear to apply to the age of the loan, which is an important explanatory variable in both APRA's PANAMA (Colman et el.) and in the Federal

Reserve's model (Calem and Follain 2004), so it is worth being more explicit about why this factor was not modelled within TUI.

The first reason is that the observed average relationship between age and default may, in part, be an artefact of a particular kind of historical experience which may not apply in the future. Most historical data reflects strong nominal house price inflation and, consequently, borrower's LVRs which fall, reasonably rapidly, from the point of origination. This in turn means that equity increases and hence the likelihood of default will fall rapidly with time. This pattern may overstate the sensitivity of default rates to time on the books in periods when house prices are relatively stable.

The second reason is that the assumption that the time profile of mortgage portfolio will be similar between banks and risk buckets is well founded when we are dealing with large banks. We have no reason to believe that a particular bank's portfolio is materially more or less seasoned than the average and that average behavioural coefficients are not picking up the appropriate level of risk.

Other technical assumptions

LVR and DSR values are current at the start of the modelling period.

Only interest payments are included in the debt servicing ratio. There is no allowance for scheduled principal payments. This assumption reflects the fact that banks readily place borrowers who come under stress on an interest only basis to reduce that pressure.

For a given DSR borrowers are heterogeneous with respect to their capacity to withstand financial shocks. This unobserved (within the model) heterogeneity is due to first, differences in other resources that can be used to pay the mortgage. This might include, for example non-housing assets, and the likelihood and amount of wider family assistance. Second, there are differences in the 'objective' level of stress imposed by a given DSR. Some families might find a forty percent ratio reasonably comfortable, while for others it could mean that they are right be on the margin of what they can afford. Third, there are differences in borrowers' inclination to fight to hold on to their residence.

This heterogeneity means that only a proportion of the borrowers in a particular risk bucket will come under stress for any particular change in the levels of unemployment and interest rates.

Each risk bucket consists of a large number of individual loans such that the number of loans that are subject to stress is always 2000. This stressed sale figure was selected to ensure that the impact of idiosyncratic housing draws on portfolio default rates and loss rates was reasonably stable.

The computation process

There are six stages in the computation process.

1. The mean expectations for the macro variables over the three year model horizon are set. In the benchmark run it is assumed that there is no expected change in interest rates and unemployment over the simulation horizon but the house price index will increase by seven and a half percent.

2. The macro-engine uses the Monte Carlo technique to draw end of period values for the house price index, the mortgage interest rate and the unemployment rate from the specified joint distribution of these variables. The joint distribution is defined by the normality assumption and the particular inputs of the volatilities and correlations of the three macro-economic variables. Each draw of the three variables represents a particular macro-economic scenario. This process is repeated 10,000 times to generate an array of 10,000 scenarios.

3. For each macro-scenario a behavioural equation determines the probability of a stressed sale (PrSS) for each risk bucket.

4. For each macro-economic scenario 2000 draws are made from a distribution of individual house prices. These draws are added to the house index value to generate a set of 2,000 individual house price changes.

5. For each of these 2,000 house price changes the model establishes whether the value of the loan is greater than the net proceeds if the house were sold. If it is then the loan defaults. The model then calculates the probability of default as the average of the 2,000 individual house events.

6. For each default event the model calculates a loss given default. When a loan defaults the value of the property is further reduced because houses subject to foreclosure attract less than the 'normal' market price. A further adjustment is made for delays in realisation process which reduces the present value of the eventual return from the sale of the house. The average of the 2,000 losses is the LGD for the risk bucket.

From the distribution of loss outcomes the model then calculates capital risk weights and other summary risk indicators. The process generates outputs for a specified risk bucket. Outputs for a given portfolio are calculated by combining the separately estimated outputs of all of the portfolio's risk buckets and then weighting them by their share in the bank's portfolio. Table 1 sets out a numerical example of the computation process. The base run coefficients, which are described in section 3, were used to generate the outputs. The first three entries in each column describe the macro scenario and the next three the model outputs that generate the scenario's loss rate.

Table 1
Model computation example

Risk bucket with 70 percent LVR, 30 percent DSR 3 year outputs				
Scenario inputs and outputs	*Scenario 1*	*Scenario 2*	*Scenarios 3-10,000*	*Risk bucket Average (average of scenarios 1-10,000)*
Interest rate change	+2 percentage points	+1 percentage point	..…….	0
Unemployment change	+3 percentage points	-1 percentage point	………	0
House price index change	-15 percent	+10 percent	……….	+7.5 percent
Probability of stressed sale	8.65 percent	3.22 percent	……….	4.46 percent
Probability of default	1.48 percent	0.10 percent	……….	0.64 percent
Loss given default	37.9 percent	37.09 percent	………..	36.90 percent
Loss rate	0.56 percent	.03 percent	………..	0.24 percent

2.3 Formal Specifications

Probability of a Stressed sale

The probability of a stressed sale is specified as follows:

$$Pr\ SS = \beta_{DS} * D + \beta_1 * \Delta DSR^{\gamma} + \beta_{DS} * [\beta_3 u + \beta_4(\Delta u)^{\alpha}] \quad (1)$$

Where:

D is a constant that captures demographic factors.

β_1 captures the linear element of the sensitivity of the probability of a stressed sale to changes in mortgage interest rates. It is set to zero if DSR falls.

γ is an exponent that recognises the non-linear impact of a change in the debt servicing ratio.

β_3 is the coefficient of default with respect to the opening unemployment rate u.

β_4 is the coefficient of default with respect to the change in the unemployment rate over the model period. Note that β_4 is set to zero if unemployment falls.

α captures the non-linear effect of an increase in unemployment.

$$\beta_{DS} = \begin{cases} 0 & \text{if } DSR < T \\ 1 & \text{if } DSR > k \\ (DSR - T)/(k-T) & \text{otherwise} \end{cases}$$

T is a minimum that (final state) DSR must exceed before it has any effect on a borrowers mortgage servicing capacity. It is set at 10 percent.

k is the DSR level where unemployment and demographic effects have their full impact.

This specification is easier to understand when it is broken into its parts.

Demographic component

$$\beta_{DS} * D \tag{2}$$

D, the demographic coefficient, and as noted in the introduction, captures a range of idiosyncratic variables which could cause a stressed sale. It is expressed as the proportion of the portfolio that will have to make a stressed sale over the 3 year time period.

It is assumed that the incidence of these idiosyncratic events is not affected by interest rate and unemployment changes. This is a simplification as it is likely that there will be some interactive effect. For example, an increase in interest rates could put family budgets under pressure and lead to an increase in marriage split ups. However, given the lack of any evidence on the size of such an effect, and their likely small effect relative to the model's overall outputs, this interactive effect has been ignored.

The term β_{DS} captures the fact that the probability of a stressed sale from demographic factors will not be constant across all DSR risk buckets. Financial mis-management, for example, is more likely to result in a stressed sale when the household's DSR is 40 percent than when it is 10 per cent. It is assumed that there is no demographic effect below a threshold level which has been set at a DSR of 10 percent. The effect then rises linearly until the full demographic effect is felt when the DSR is 25 percent.

The demographic factor is relatively undeveloped in this model and the effect of all the components are captured by a single coefficient. This is adequate for estimating capital because the demographic coefficient has only a small impact on the capital requirement. With stress testing its relative importance will depend on the size of the test event. With large, relatively infrequent events, the demographic factor also has a relatively minor effect. With high frequency events, such as the bottom of a normal credit cycle it could be a dominating influence. For this sort of analysis a more developed demographic module would be useful.

Debt Servicing Ratio

$$\beta_1 * \Delta DSR^{\gamma} \tag{3}$$

This term says that stressed sales are a function of changes in the debt servicing ratio. Note that only changes in the DSR affect the stressed sale rate in this part of the equation. While high DSR levels are in themselves positively related to the probability of a stressed sale, this effect is captured by the structure which determines the impact of the demographic variable. The only variable impacting on the DSR is the interest rate. The coefficients β_1 and γ capture the idea that the probability of a stressed sale is an increasing and non-linear function of the size of the debt servicing shock.

The values of β_1 and γ reflect the unobserved heterogeneity that is not captured directly by the model. This heterogeneity affects the position and slope of the curve that defines the relationship between DSR changes and the level of stressed sales.

Unemployment

$$\beta_{DS} * [\beta_3 u + \beta_4(\Delta u)^{\alpha}] \qquad (4)$$

The first term in the brackets $\beta_3 u$ says that the probability of a stressed sale is a function of the starting level of unemployment. Even if there is no increase in unemployment over the model horizon there is still a chance of the borrower becoming unemployed because there is movement into, and out of, a given stock of unemployment. This effect is relatively small because the home owning group are different in character from the non home owners. They tend to be older and more financially secure and, therefore, have a lower chance of becoming unemployed over the model horizon.

The second term $\beta_4(\Delta u)^{\alpha}$ says that the probability of a stressed sale is an increasing non-linear function of the increase in unemployment over the model horizon. A fall in unemployment has no impact on the probability of a stressed sale.

The term outside the brackets, $\beta_{DS,}$ adjusts the impact of the unemployment terms, depending on the size of the DSR. This is the same adjustment that is applied to the demographic impact.

Economic default

If there is a stressed sale and if the net value of the individual house, after disposable costs at the end of the period, is less than the value of the loan then the borrower will default.

$$H - C < L \qquad (5)$$

There is no default if $H - C$ is greater than L

Where H = the market value of a particular house at the end of the period
C = the cost of disposing of the house

L = the loan value

Loss given default

Loss given default is defined as follows:

$$LGD = L + C - (1-\delta)H_F/(1+i+r)^{T} \quad (6)$$

Where:

- L = loan value
- C = transaction cost associated with foreclosing on a mortgage
- δ = the discount applying to a mortgagee sale
- H_F = the final sale price of a particular house
- i = the final period interest rate
- r = the risk premium
- T = the time to sell

The third term captures the time value cost of the delay in liquidating the loan. T represents the time to sell in years and r is a risk premium appropriate to a foreclosed asset. It is expressed as a margin over the normal housing lending rate than as a margin over the return on the safe asset. There is a discount δ on the value of the house because foreclosure depresses prices below the normal willing seller/ willing buyer price.

Calculation of the loss distribution, risk weights and other outputs

The 10,000 loss outcomes are set out as an array by frequency and by the amount of the loss as shown in figure 2. The Basel II IRB risk weight regime requires capital to be held for unexpected losses which is the difference between the expected loss and the loss at the prescribed confidence level. This is the amount UL. TUI generates a figure for expected loss, which is the average loss over the 10,000 scenarios, and the loss at the specified confidence level, to obtain a required level of capital.

Figure 2
Derivation of capital requirement

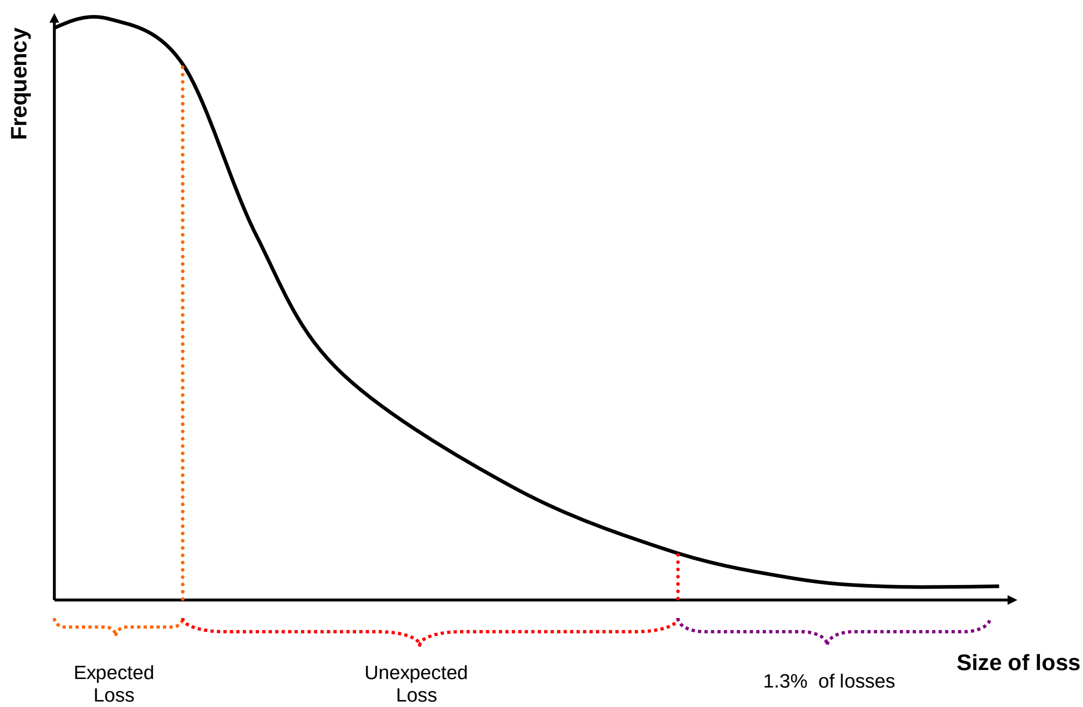

For a given required level of capital K (expressed as a proportion of the loan) the risk weight is calculated using the formula:

$$RW = K \times 12.5$$

A capital requirement for a risk bucket of eight percent translates to a risk weight of 100 percent.

Stress Test Mode

TUI can also be used in a stress test mode. It will generate expected losses together with the associated PD and LGD for a three year period given a particular set of macro inputs. This is done by setting the volatility figures at zero and setting the expected outcomes inputs at the desired scenario values. If the scenario is longer than three years then a succession of three year outputs can be generated.

The required scenario values are the change in the housing price index, the change in the interest rate, and the change in the unemployment rate. It is also possible to change any of the base model behavioural coefficients.

3 Model inputs

While one of TUI's strengths is that it allows the user to input their own estimates of a range of coefficients that describe the structure of the macro economy, borrowers' behaviour and the micro-structure of the housing market, it is important to have a reasonably robust and credible set of inputs that define a benchmark scenario.

In this section we set out a set of benchmark inputs and describe the economic logic, sources and judgements behind them. We have not always achieved good or acceptable standards with all of the inputs and table 2, which summarises the robustness of the inputs, shows some obvious shortcomings.

That being said, we do not believe that difficulties with data inputs are a compelling argument against the value of the TUI modelling exercise.

- Some key outputs, for example, the shape of the loss function, are relatively robust to the magnitude of particular inputs.
- It has been possible to find indirect checks on the quality of some of the outputs.
- This set of inputs should be regarded as provisional inputs into a prototype model. It is hoped that they will generate ideas about better ways of modelling the failure process; the collection of better data, and better ways of depicting the structure of future macroeconomic events.
- Banks have to somehow make judgements about the risk of mortgage loans regardless of the extent of uncertainty as to the true state of the world. The point, therefore, is not whether this model meets some absolute standard of accuracy or robustness but whether it is better than alternative ways of calculating bank capital or provides a useful supplement to them.

3.1 Macroeconomic inputs

The benchmark macroeconomic inputs represent the consensus view of a group of Reserve Bank of New Zealand economists.

Expected values

The expected value of the change in interest rates and unemployment is normally set at zero. House prices are expected to increase by 7.5 percent over the three year model horizon which represents a small real increase.

Volatilities

House price volatility

The benchmark model standard deviation of 17.5 percent was based on a consideration of New Zealand and international evidence. Our empirical estimate of the standard deviation of real house prices, over a three year horizon for New Zealand over the period 1974 - 2005, was 12.5 percent. With respect to foreign markets we relied on the IMF (2004) report on banking crises. It showed that major house price crashes were relatively frequent, with twenty falls in real house prices of over fourteen percent in fourteen OECD countries over a period of 20 years. The crashes lasted for an average of four years and the average fall in real house prices was twenty-eight percent. While our New Zealand data seemed to suggest that the New Zealand housing market is more stable than the OECD markets we have placed more emphasis on our prior that New Zealand is more likely to be an average performer and that our lower measured figure is an artefact of the particular observation period. Our 17.5 percent standard deviation estimate also has regard to what appears to be evidence of a fat tail in the international housing price distribution data.

Mortgage interest rate volatility

The base run estimate of 2.5 percent had regard to the following evidence:

- The three year standard deviation of New Zealand real interest rates for the period 1985-2005 was around 2.25 percent.
- The average increase over the three years of the UK housing slump, which we have used to calibrate TUI's sensitivity to changes in the debt servicing ratio, was around 3.5 percentage points. If this slump was a two standard deviation event, then this implies a standard deviation of just under 2 percentage points.

A standard deviation of 2.5 percentage points was selected because:

- New Zealand interest rates can be expected to be more volatile than UK rates.
- An upwards adjustment to the measured New Zealand real rate volatility figure to account for the impact of inflation volatility on nominal interest rate volatility is required.

While a standard deviation of 2.5 percentage points might seem on the low side it should be noted that this figure is an average over three years and not a measure of point to point volatility. The former measure is more stable than the latter.

Unemployment

The assumption of a standard deviation of 3 percentage points does not seem unreasonably large given the kinds of employment shocks experienced by New Zealand and similar countries over the last century. A case could be made for a higher figure but the model's capital requirement output is not very sensitive to this variable.

Correlation coefficients

(i) House price/interest rate
The house price/interest rate 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 together. Conversely, a shock to the New Zealand dollar risk premium would put upward pressure on interest rates and downward pressure on house prices. These competing influences are reflected in the New Zealand empirical evidence which shows cross-correlation fluctuating on either side of zero depending on the sample period.

A key consideration is that it is 'tail' correlation rather than the average correlation that matters when modelling the kind of events that generate capital requirements. Here some of the evidence points to a negative correlation. The IMF study cited above showed that house price crashes were typically immediately preceded by interest rate increases.

It is more likely that the downturn or stressed interest rate house price cross-correlation in a small open economy such as New Zealand will be negative. In a large country like the US, monetary policy will typically only tighten in the face of falling house prices if there is a genuine risk of inflation increasing. In New Zealand, on the other hand, a negative domestic shock, which taken by itself would have negative implications for inflation, could be associated with an increase in the risk premium on the New Zealand dollar and a need for the Reserve Bank to increase the cash rate to support its value. Similarly, a generalised 'flight to quality' in response to a global shock could also see New Zealand interest rates rise and US interest rates fall as the New Zealand economy and house prices turned down.

On the other hand it is obviously possible to point to international episodes where large house price declines have not coincided with higher interest rates and where interest rates have fallen as the economy and house prices have weakened. We settled on a value of -0.3 which gives a greater weight to the negative correlation drivers but still provides a place for the conventional demand side positive correlation between house prices and interest rates.

(ii) House price/unemployment
The base run coefficient of -0.5 is consistent with the actual 1984 - 2004 New Zealand experience. As this variable is not critical to the model's output and because the -0.5 figure seems economically plausible it was not subject to intensive scrutiny.

(iii) Unemployment/interest rates
The economics here are similar to those underpinning the house price/unemployment estimate so the same coefficient value (with the opposite sign) has been used.

3.2 Behavioural inputs

The UK housing loan loss experience in the early 1990s was used to generate a measure of the relationship between the probability of a stressed sale and unemployment and mortgage interest rates. The British experience was used because:

- The structure of the mortgage market, in particular the absence of long term fixed rate loans, is closer to the New Zealand market than, for example, the US market.
- It combines a major interest rate shock with a large negative house price shock.
- Statistical information and analysis is more readily available than for housing shocks in other countries.

Some key, stylised, facts about the British experience are as follows:

- The average increase in residential mortgage interest rates over a period of three years was around 3.5 percentage points, or over 50 percent of the pre-shock interest rate.
- House prices fell by around 30 percent.
- Average unemployment over three years (allowing for a lag of one year) increased by 60 percent from the pre-shock level.
- The stock of mortgage arrears of six months or more increased by a maximum of 600 percent from the pre-shock level.
- Annual loan write-offs by the major building societies increased by between one hundred and two hundred times their pre-shock levels.

Sensitivity of stressed sales to interest rates

There are four main steps in the derivation of the behavioural coefficients for the sensitivity of stressed sales to increases in the debt servicing ratio caused by interest rate increases. They are:

1 Derive a figure for the aggregate rate of stressed sales over three years of the most acute period of the UK housing shock.

2 Allocate the aggregate stressed sales figure between those caused by the increase in unemployment and those cause by the interest rate driven increase in the debt service ratio.

3 Calculate the implied overall stressed sale/interest rate sensitivity coefficient

4 Calculate the interest rate sensitivity coefficient for each DSR tranche.

Step One: Generate an aggregate stressed sales rate.

While there has been considerable analysis of the UK housing loan default experience since the crisis, little of this was directly useful in estimating a three year stressed sale rate. The main reason is that most of the work has been concerned with explaining changes in the stock of defaulted loans (probably because this data was readily available). TUI, however, needs data on the flow of stressed sales over a period of interest.

Our approach was to 'reverse engineer' an estimate of stressed sales from housing loan write-off figures for the major UK Building Societies obtained from the Bank of England. This involved the following steps:

- Estimate the three year loss rate from write-off data. This was 1.7 percent.
- Calculate the value of loans in default using an assumed average LGD of 35 percent.
- Adjust for the proportion of stressed sales not leading to economic default.

This generated an aggregate stressed sale figure by volume of lending of 7.3 percent.

Step two: Allocate aggregate stressed sales to DSR and to other causes.

Here we have relied on the work by Whitely, Windram and Cox (2004) on the determinants of residential mortgage borrowers falling into arrears. The analysis was conducted over the period 1988 - 2002 but was largely driven by the housing shock of the early 1990's. We have assumed that their results, as to the relative importance of the main risk drivers, can be applied to the shorter, more intense, period of the crisis that we are interested in. The key result was that around 80 percent of defaults are explained by interest rates movements and most of the remainder by changes in unemployment.

The 80/20 split between the interest rate and unemployment effect is intuitively appealing. The reason for the preponderance of the interest rate effect is that an interest rate increase impacts on all borrowers whereas unemployment only impacts on those who become unemployed. So while the probability that a borrower who becomes unemployed will become stressed is much higher than the borrower who is subject to an interest rate increase, the much larger numbers affected in the latter case drives the overall impact.

Step three: Calculate the aggregate interest rate sensitivity.

A 50 percent increase in the debt servicing ratio increases stressed sales by 6 percentage points.

Step four: Calculate the sensitivity of stressed sales to interest rates changes for DSR tranches.

The challenge here is that we only had an aggregate stressed sale figure, no information on how UK portfolios were allocated between DSR buckets and only weak information on the sensitivity of default by DSR tranches to increases in the debt servicing ratio. Our strong theoretical and common sense prior, however, was that higher DSR buckets were more sensitive to increases in servicing costs than lower ones.

We have support for this proposition from an analysis of default determinants by May and Tuleda (2005) who used longitudinal household data. They found that borrowers' debt servicing ratios helped explain mortgage defaults for ratios of over 20 percent but was not significant for servicing ratios under that. The lack of data prevented them from exploring the full functional form but we have used their result in our calibration of the SS sensitivity coefficient for our lower DSR tranches.

Absent information on the structure of UK portfolios we assumed that the distribution by DSR tranche, immediately prior to the shock, was the same as the New Zealand average portfolio in 2004 generated by the Statistics Department's Household Expenditure Survey.

Given the overall sensitivity of stressed sales to the interest rate increase generated in step three, a stressed sale/interest rate change functional relationship across DSR risk buckets was constructed. This had to meet the consistency requirement that their weighted average had to sum to the aggregate sensitivity.

Figure 4 shows the relationship between DSR and SS for the 50 percent debt servicing shock of the British housing crisis. The DSR figures on the horizontal axis shows the debt servicing ratio before the shock. The vertical axis shows the probability of a stressed sale.

The SS/DSR line is weakly non-linear to demonstrate the capacity of the model to handle a non-linear relationship. The non-linearity factor was kept small because we have no information or priors about the degree of non-linearity in the relationship although we believed the shape of the curve is convex. Thus our debt servicing sensitivity coefficients are, weakly, anchored at the lower end but not at the higher end and a range of plausible formulations could have met the aggregate sensitivity test. The relative sensitivity of the DSR risk buckets is, then, a weakness in the model when it is used to calculate the relative risk of DSR buckets. It is less of a problem when the problem is to estimate the risk of a total portfolio because of the overall sensitivity constraint.

Subsequent to the building and testing of TUI we became aware of a set of default rates by DSR published by Fitch (2007). This appears to show that defaults rates are significantly less sensitive to the debt servicing ratio than the TUI calibration suggests. The difference in the default rate of the 20-25 percent DSR risk bucket and the highest DSR bucket is about 110 percent for TUI but only around 60 percent for Fitch. There may be a technical reason for some of the difference. The Fitch data appears, on average, to reflect relatively benign conditions than those TUI is calibrated to. That is there is a greater influence of idiosyncratic default events which probably generate a less sensitive relationship between default rates and DSRs than large systemic interest rate shocks.

However, the Fitch data suggests some further caution in the interpretation of our DSR default rates sensitivities and that some reduction in the relativities could be appropriate.

Figure 2
Relationship between DSR and stressed sales rate

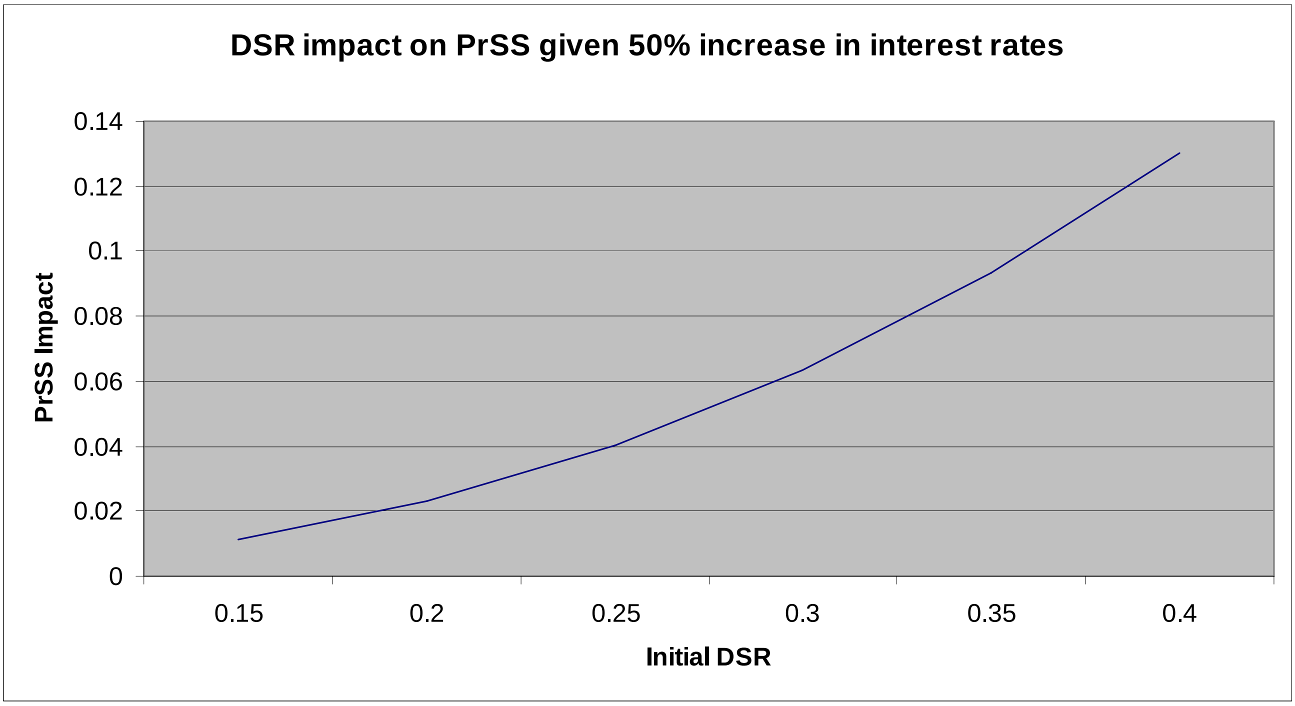

Sensitivity of stressed sale rate to changes in unemployment and the level of unemployment

The impact of the level of unemployment on stressed sales was based on recent default data from a New Zealand bank.

Our analysis of the effect of changes in unemployment on stressed sales is not strong. While the Whitely et al (2004) study gave us some information on the relationship between aggregate unemployment and defaults in the UK this did not directly provide information on changes in the unemployment of wage and salary earners and defaults by that group The unemployment variable will be strongly correlated wth a downturn in self-employed incomes and it is likely that a material part of the defaults came from the latter group. Absent any better information, the coefficient for the proportion of unemployed who move to a stressed sale was set at 0.7. This figure looks very high, and might be, but it should be remembered that this represents a three year effect. Further, our replication of the FSAP stress test, which is reported below, suggests that the banks' losses implied a higher stressed sale/unemployment coefficient.

Demographic coefficient

New Zealand banks' IRB modelling has provided some useful guidance as to the likely size of the demographic coefficient. They have been generating default rates of under two percent over a three year horizon. Nearly all of these defaults would have been generated by demographic events rather than by increased unemployment or by increasing interest rates. The IRB definition of default is similar in some respects to the TUI stressed sale concept with the following differences: It excludes stressed sales that do not result in the borrower being 90 days overdue, and it includes overdues that are subsequently cured that are not captured in the stressed sale concept. If it is assumed that these components are of the same magnitude then the demographic stressed sales rate can be assumed to be equal to recent Basel II IRB defaults. We have therefore assumed a demographic stressed sale rate of two percent.

3.3 Idiosyncratic housing risk volatility

A distribution of conditional individual house price changes was derived by deducting the general house price index change from the price changes of sets of paired house sale prices obtained from Valuation NZ. The data showed that conditional volatility of individual house prices was high, was positively skewed and that both the positive and negative sides of the distribution had fat tails. To some extent the positive skewness would have been an artefact of a data period when there was generally a strong upwards trend in nominal house prices and it would have also been affected by factors such as renovations and non-economic transfers for tax and family reasons. This positive skewness could not be relied on in a stress situation so it was assumed that the overall distribution was normally distributed with a standard deviation of 10 percent. However, the information in the fat tail of the downside of the distribution was retained. A curve was fitted to this tail distribution and spliced to the truncated normal distribution to derive the overall individual house price distribution.

3.4 Loan Loss Coefficients

Transaction costs

The cost of realising the security after a foreclosure was assumed to be 5 percent of the value of the house. Real estate sales fees costs constitute the greater part of this cost.

Discount rate

The discount rate on foreclosed loans is 5 percent above the mortgage rate. This represents a risk premium of about six percentage points above the risk free rate which is an appropriate rate for foreclosed debt.

The discount rate for insurance receivables is the lending rate to the insurance company. This will vary depending on the circumstances of the company but most weight should be given to the rate in a stress situation. We have assumed a risk premium of 200 basis points.

Time to collection

Time to collection of a defaulted loan is set at 1.25 years based on some fairly anecdotal evidence of the average time to recover loans in a stress situation. A single figure has been used while it might have been more appropriate to use a distribution of possible time lags. This could be gathered from banks' actual experiences. Using a distribution will not make any appreciable different to the higher and medium risk buckets but it will probably increase the risk of the lowest risk buckets by a moderate amount.

Discount for foreclosure

We had regard to US studies cited by Pennington Cross (2003) that estimated a discount for foreclosure in the order of 23-24 percent. We were also informed, by a New Zealand bank, that they used, as a rule of thumb, a discount figure of 15 percent. We believe that the US figures would be too high for our model because we are probably already picking up some of the measured discount in the model's downside idiosyncratic price outcomes mentioned above. We settled, therefore, on the New Zealand bank estimate of 15 percent. We have also used a point estimate for this variable when, a distribution would have been more appropriate. The use of a distribution might, again, increase the risk of the lower risk buckets but would not have much, if any, impact on the medium and higher risk buckets. It will be a simple matter to build a distribution into the model although setting the parameters of the model will be largely a matter of guess-work.

Figure 3
TUI coefficient input page

TUI Credit Loss Distribution Model

Model Run Number	1
Model Run Description	Base

Run Tui

1 Base Model Inputs

Macro economic

	Interest rate	House price index	Unemployment	Min U	4
Mean	0	7,500	0		
SD	2.5	17,500	3		
Initial Values	7	100,000	4		

Correlation matrix

	i	h	u	y
i	1			
h	-0.3	1		
u	0.3	-0.5	1	
y	0.5	0.5	-0.5	1

Demographic Coefficient	0.02

Behavioural Coefficients		*Loss parameters*	
DSR Linear coeff	0.023	Foreclosure discount	0.85
DSR non-linear coeff	2.5	Risk premium percentage points	5
Unemployment coeff	0.06	Time to sale - years	1.25
Delta Unemployment coeff	0.7	Transaction costs	5%
Delta U DSR Threshold	0.1		
DSR/delta U scalar	0.15		
Delta U non-linear coeff	1		

Iterations number	10,000
Percentile	0.987

2 Insurance Module

OFF

Contract Type	AUSSIE
LVR threshold	0.8
Payout rate	0.8
Failure rate	0

3 Fixed interest Module

OFF

	Shock Response	Portfolio Weight
1 YR	20%	33%
2 YR	60%	33%
3 YR	20%	33%

3 Self Employed Module

OFF

Number	2,000
Indiv income s.d.	1
Prop Self-employed	0.1
Mean Scalar	0.2

4 Rational forebearance Module

OFF

Proportion of forbearance	0.20
Proportion of loss saved	0.20

Table 2 provides a summary of coefficient values; data sources; a 'traffic light' assessment of the robustness of our estimates; and a measure of the relative importance of each variable.

Table 2
Summary of model inputs

Variable	Value	Sources	Quality	Importance
House price volatility	17.5 percent	NZ empirical IMF		High
Interest rate volatility	2.5 percentage points	NZ UK		High
Unemployment volatility	3 percentage points	NZ UK		Medium
House/interest rate corr.	-0.3	Judgement		High
House/unemployment corr.	0.5	NZ empirics		Low
Interest rate/ unemployment corr.	0.3	Judgement		Low
Demographic	0.02	NZ banks		Low
Idiosyncratic house prices	Distribution	VNZ Judgement		Medium
SS interest rate sensitivity	Distribution	UK data and studies NZ data		High
SS unemployment sensitivity	0.0024	Judgement		Low
SS unemployment change sensitivity	0.7	Judgement		Medium
DSR risk bucket sensitivity to default	NA	Judgement		Medium
Discount rates	5 percentage points	Judgement Markets		Medium
Foreclosure discount	15 percent	US literature NZ bank		Medium
Transaction costs	5 percent	NZ banks		Low
Time to collection	1.25 years	NZ banks Judgement		Low

4 Base Model Results

This section presents a full array of model outputs for the base run.

4.1 The loss distribution

Figure 4 sets out the respective loan loss distributions for a high (40 percent DSR, 100 percent LVR), a moderate (30/70), and a low (15/60) risk buckets and for an illustrative loan portfolio whose portfolio weights are derived from 2004 household expenditure survey. The vertical axis shows the percentage of the portfolio and the horizontal axis shows the average default rate. The graph shows that the low risk bucket loss distributions are highly concentrated to the left and that there is a shift to the right as risk increases.

Figure 4
Loss distribution by risk bucket and full portfolio

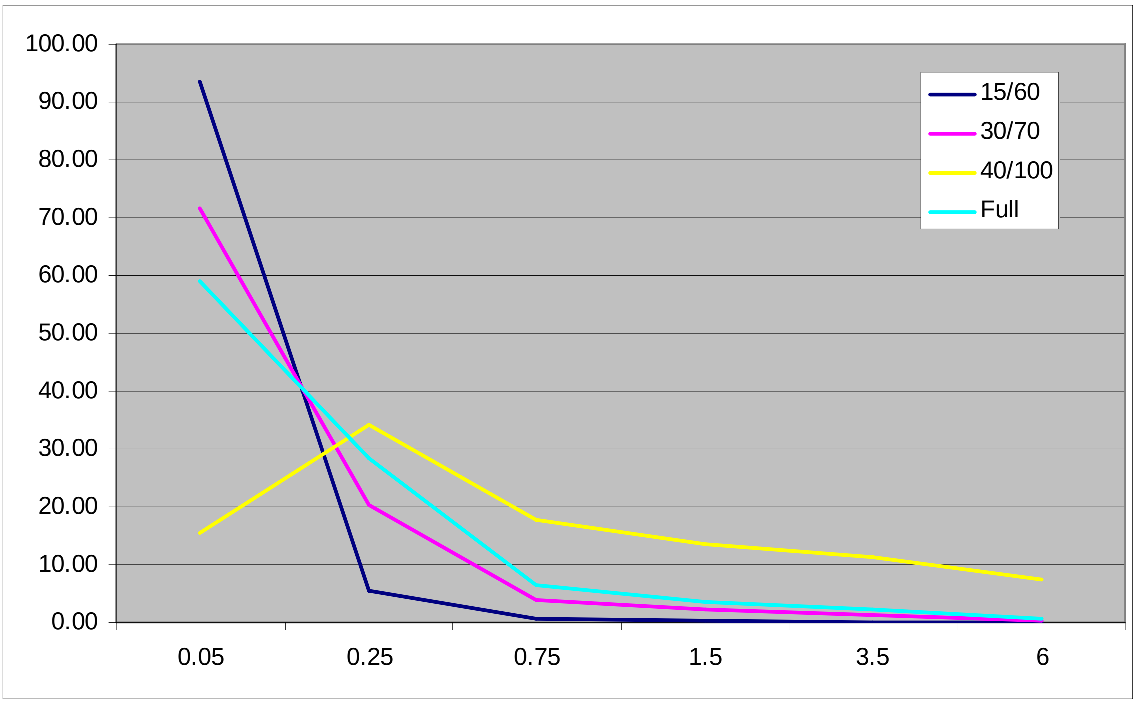

4.2 Risk weights by Risk Bucket

Table 4 shows one of TUI's key outputs - risk weights by risk buckets. The information is also expressed, equivalently, as a total (tier one and tier two) capital ratio in table 5. Through the rest of the paper we have expressed risk in risk weighted terms. These numbers can be readily converted to capital ratios by dividing the risk weight by 12.5.

The most obvious features of the structure of risk weights are, first, that risk strongly increases with the LVR. Second, risk also increases with the average debt servicing ratio although as noted above, this is largely an artefact of the particular assumptions we have imposed with respect to this sensitivity.

Third, risk is material for the highest LVR and DSR risk buckets and conversely very low for the low LVR/DSR buckets. The results show that the old bankers' rules about debt servicing and leverage (limits of 30 and 70 percent respectively were typical), which were designed to exclude risky loans from the portfolio, were pretty much on the mark.

Table 4
Risk weights by risk bucket base case

LVR / ***DSR***	***100***	***90***	***80***	***70***	***60***	***50***
40	*188*	*143*	*88*	*46*	*18*	*6*
30	*101*	*79*	*52*	*28*	*11*	*4*
20	*52*	*45*	*33*	*18*	*6*	*2*
15	*37*	*31*	*23*	*12*	*4*	*1*

Table 5
Total capital by risk bucket base case

LVR / ***DSR***	***100***	***90***	***80***	***70***	***60***	***50***
40	*15.0*	*11.4*	*7.0*	*3.7*	*1.4*	*0.5*
30	*8.1*	*6.3*	*4.2*	*2.2*	*0.9*	*0.3*
20	*4.2*	*3.6*	*2.6*	*1.4*	*0.5*	*0.2*
15	*3.0*	*2.5*	*1.8*	*1.0*	*0.3*	*0.1*

4.3 Average Probability of Stressed sale by DSR

Table 6 shows the probability of a stressed sale by risk bucket. Note that because stressed sales are not driven by the loan to valuation ratio. The probability of a stressed sale is the same for each LVR tranche.

Table 6
Average probability of a stressed sale base case

DSR	***%***
40	*6.5*
30	*4.5*
20	*2.6*
15	*1.5*

4.4 Average PD by risk buckets

Table 7 shows the average probability of default by risk bucket. Note that the figures presented here are three year default rates. To put the data in the more familiar annual average terms they should be divided by three. Note also that these figures cannot be compared directly with the IRB default probabilities because the IRB definition included liquidity events where there is no economic loss. The structure of PDs by risk bucket is similar to the structure of risk weights. Risky loans have much higher default rates than traditional conservative loans.

These results point to one of the pit-falls of using historic data to draw conclusions about current expected default rates when the risk structure of a loan portfolio has changed over time. Suppose that a bank had for a long period of time implemented a relatively conservative lending policy and that its long run average default rate was approximated by the 30/60 risk bucket. Then assume that in recent years the bank relaxes its LVR and DSR limits during a benign loss experience period so that thirty percent of the portfolio is in the higher risk loans which are approximated by the characteristics of the 40/90 bucket. Because the bank is operating in a benign period it would not have experienced an increase in its loss rate as the risk profile of its portfolio changed and, based on its historic portfolio loss experience, might expect that its long run default rate was still 0.22 percent. Its true loss rate, however, will have increased to 0.97 percent (0.22 x 0.7 + 2.71 x 0.3).

Table 7
Average three year PD (%) by risk bucket

LVR / ***DSR***	***100***	***90***	***80***	***70***	***60***	***50***
40	*3.94*	*2.71*	*1.66*	*0.87*	*0.35*	*0.13*
30	*2.66*	*1.84*	*1.08*	*0.53*	*0.22*	*0.09*
20	*1.59*	*1.11*	*0.65*	*0.33*	*0.13*	*0.05*
15	*0.96*	*0.67*	*0.40*	*0.20*	*0.09*	*0.03*

4.5 Average LGD and stressed LGD by risk bucket

Table 8 shows average and stressed LGDs by LVR tranche. By assumption there is no relationship between LGD and DSR.

The stressed LGDs were calculated by taking the average of the 11 LGDs centred on the 0.987 percentile. This is intended to measure the exact event that defines the capital requirement for each risk bucket but an average is taken to smooth out random sampling effects. The key feature of the results is the lack of sensitivity of both the average and stressed LGDs to the risk bucket LVRs. There is almost no relationship at all with average LGDs and with stressed LGDs they are only moderately higher with the higher LVRs. While these patterns do not seem intuitive they do follow logically from the model structure which assumes a large fixed cost due to time value and the foreclosure discount when any loan goes into default (for a more detailed discussion of the drivers of this result see appendix A). Our results are also supported by empirical evidence. Australian mortgage insurance loss rate data also shows little relationship between loss rates and LVRs.

Table 8 Average and stressed LGD (%) by risk bucket

LVR	*100*	*90*	*80*	*70*	*60*	*50*
Average LGD	*39.0*	*38.2*	*37.7*	*37.7*	*38.7*	*40.3*
Stressed LGD	*46.4*	*44.1*	*41.6*	*39.7*	*39.0*	*40.2*

The relatively small difference between the average, and stressed LGD, distributions of LGDs is, in part, an artefact of the LGD calculation methodology. Two important drivers of LGD, the time to recovery of foreclosed loans and the discount on the sale price for a bank foreclosure are modelled as single expected values rather than as distributions. In reality there will be a spread around the average figure with both the stressed time to disposal and the foreclosure discount being higher than the average. Our point estimates for both figures were calibrated to downturn conditions so the stressed LGD, which feeds into the model's capital outputs, is reflective of our best judgements. LGDs for more moderate outcomes will, on the other hand, tend to be overstated. This will mean that TUI might tend to overstate losses in moderate shocks when it is used in stress test mode.

4.6 Impact of idiosyncratic housing risk

Table 9 shows the set risk weights calculated without the influence of idiosyncratic housing risk. A comparison with the base run outputs in table 4 provides some insights into the impact of idiosyncratic house price movements. For the higher risk buckets, idiosyncratic house price risk actually reduces the risk weight. In the lower risk buckets, however, risk weights are materially increased and the zone where capital is required is extended to significantly lower LVR tranches.

The result for the lower risk buckets was expected but the reduction for the high risk buckets came as a surprise. Our explanation for this phenomenon is that the idiosyncratic risk introduces a diversification effect compared to the case where all house prices are driven by a single index. In the high risk buckets, where systemic risk is strong, this diversification effect outweighs the additional volatility that idiosyncratic risk introduces into the portfolio. In the low risk buckets where there is less, or no systemic effect, there is less or no diversification effect and the effect of introducing additional volatility dominates to impart an overall increase in risk.

Table 9 Risk weights without idiosyncratic housing volatility

LVR / *DSR*	*100*	*90*	*80*	*70*	*60*	*50*
40	*196*	*160*	*98*	*35*	*0*	*0*
30	*105*	*89*	*59*	*28*	*0*	*0*
20	*53*	*47*	*38*	*20*	*0*	*0*
15	*38*	*33*	*26*	*12*	*0*	*0*

4.7 New Zealand portfolio results

Table 10 sets out some summary outputs for a portfolio generated from data from the Statistics Department's 2004 Household Economic Survey which should be broadly reflective of the average LVR/DSR structure of New Zealand bank portfolios at that time.

Table 10 HES portfolio summary results

	Percent
LGD	38.6
Stressed LGD	40.7
PD (3Yr)	0.7
Risk weight	24.7
Loss rate (3Yr)	0.3

5 Sensitivity Analysis

The following tables show the sensitivity of risk weights to changes in the coefficient values for the key macroeconomic and behavioural variables in the base run. The results are presented for a small selection of risk weights to give a flavour of the overall results.

5.1 Demographic Coefficient

Benchmark run coefficient value: 2 percent

The clear message from table 11 is that the demographic coefficient is not a significant driver of housing risk weights. A doubling of the coefficient increases the 30/70 risk weight by only 5 percentage points and even an increase in the coefficient to 10 percent increases the risk weight for this risk bucket by only 50 percent. The reason for this lack of sensitivity is that the underlying demographic factors are assumed to be stable and it is only the interaction with other variables that adds to unexpected loss.

Table 11 Impact of Demographic coefficient on risk weights

Coef. / *Bucket*	*2 Percent*	*0*	*3*	*4*	*10*
40/90	143	135	146	149	168
30/70	28	24	31	33	50
20/60	6	5	7	8	13
Portfolio	24.7	22.2	26.1	27.5	36.0

5.2 Macroeconomic coefficients

Interest rate volatility

Benchmark coefficient value: 2.5 percentage points
Table 12 shows that the riskier risk buckets are highly sensitive to changes in the interest rate sensitivity assumption but that medium risk buckets are relatively stable.

The contribution of interest rate volatility to overall risk becomes apparent when the volatility figure is set to zero. It is the major driver of risk for the high risk buckets but is responsible for a smaller part of total risk for the low risk bucket.

Table 12 Impact of interest rate volatility on risk weights

Vol.* / *Bucket	***2.5 percentage points***	***0***	***1.5***	***2***	***3***	***4***
40/90	*143*	*28*	*54*	*87*	*215*	*390*
30/70	*28*	*14*	*18*	*22*	*37*	*62*
20/60	*6*	*3*	*5*	*6*	*8*	*11*
Portfolio	*24.7*	*8.2*	*12.9*	*17.4*	*35.2*	*60.9*

House price index volatility

Benchmark coefficient value: 17.5 percent

Risk weights are highly sensitive to volatility in the housing index. When there is no volatility in the index the model runs off idiosyncratic housing price volatility. This has a material impact with the high risk buckets but generates low results at the middle and lower end of the risk spectrum. As house price index volatility is introduced and increased then there is a substantial effect across the risk spectrum but a proportionately larger impact on the lower risk buckets.

Table 13 Impact of housing index volatility on risk weights

Vol.* / *Bucket	***17.5 percent***	***0***	***10***	***15***	***20***	***25***
40/90	*143*	*35*	*82*	*123*	*158*	*187*
30/70	*28*	*3*	*9*	*19*	*40*	*61*
20/60	*6*	*1*	*2*	*4*	*11*	*26*
Portfolio	*24.7*	*6.2*	*13.0*	*20.2*	*30.3*	*43.1*

Unemployment volatility

Benchmark coefficient value: 3 percentage points
Unemployment has a lower impact on risk weight levels than either house price or interest rate volatility buckets. Even the large proportionate increase in the volatility assumption from 3 to 6 percent increases the average risk weight by only 20 percent.

Table 14 Impact of unemployment volatility on risk weights

Vol. / ***Bucket***	***3 percent***	***0***	***2***	***4***	***6***
40/90	*143*	*133*	*138*	*146*	*152*
30/70	*28*	*20*	*25*	*31*	*35*
20/60	*6*	*4*	*6*	*7*	*9*
Portfolio	*24.7*	*21.0*	*23.5*	*26.3*	*29.4*

Correlations

Benchmark correlation values	
Interest rate / house price:	**-0.3**
Interest rate / unemployment:	**0.3**
House price / unemployment:	**0.5**

The key driver in the correlation matrix is the correlation between the house price index and interest rates. The correlation of unemployment and interest rates is not so important but as it reflects the same underlying economic forces that drive the house price interest rate correlation the sensitivity to changes in the two correlations has been analysed jointly. For example when the house price/interest rate correlation is set at -0.5, the interest rate/unemployment correlation is set at 0.5, and when the former is set at 0.5 the latter is set at -0.5.

The results in table 15 point to the importance of the structural relationship between interest rates and house prices in stress situations. If that structure is dominated by scenarios where interest rates fall when there are falls in house prices then it will be a very benign world from a housing risk perspective. A positive correlation of 0.5 generates an average risk weight of only 6.7 percent for our benchmark portfolio and the high risk bucket risk weight is still only 30 percent.

On the other hand, a world where large house price falls tend to be accompanied by increases in the New Zealand dollar risk premium, or by monetary policy tightenings, will require materially higher capital ratios. The difference in the capital requirement between correlations of plus and minus 0.5 is around five to one.

Table 15 Interest rate / house price / unemployment correlation impact on risk weights

Corr. / *Bucket*	***-0.3***	***0.0***	***0.3***	***0.5***	***0.7***	***-0.5***	***-0.7***
40/90	*143*	*86*	*48*	*30*	*28*	*186*	*222*
30/70	*28*	*18*	*11*	*7*	*6*	*36*	*52*
20/60	*6*	*4*	*2*	*1*	*1*	*8*	*10*
Portfolio	*24.7*	*16.1*	*9.8*	*6.7*	*5.7*	*32.5*	*42.3*

Table 16 shows that there is limited sensitivity to the unemployment/house price correlation, particularly in the high risk bucket.

Table 16 Unemployment / house price correlation impact on risk weights

Corr. / *Bucket*	***0.5***	***0***	***-0.3***	***-0.7***
40/90	*143*	*136*	*141*	*144*
30/70	*28*	*24*	*26*	*31*
20/60	*6*	*5*	*6*	*7*
Portfolio	*24.7*	*23.0*	*23.9*	*25.7*

5.3 Behavioural Coefficients

The material behavioural coefficients in the model are the relationship between the probability of a stressed sale and the DSR (the debt servicing sensitivity coefficient) and the discount for default on the sale of foreclosed properties.

DSR / Pr. SS sensitivity benchmark coefficient value: .023

Note that the coefficient values in the table below are expressed as a percentage of the base run coefficient not in absolute terms because the absolute values do not have the more intuitively obvious meaning that some of the other model coefficients have.

Table 17 shows that risk weights are not sensitive to changes in the debt servicing sensitivity coefficient in the medium to lower risk buckets, although there is a stronger relationship in the highest risk buckets.

Table 17 SS/DSR coefficient impact on risk weights

Bucket \ Coeff. %	*100*	*50*	*70*	*90*	*110*	*150*
40/90	*143*	*80*	*105*	*130*	*156*	*203*
30/70	*28*	*21*	*24*	*27*	*29*	*35*
20/60	*6*	*6*	*6*	*6*	*6*	*7*
Portfolio	*24.7*	*16.6*	*19.8*	*23.1*	*26.5*	*33.3*

(ii) Foreclosure discount

Base coefficient value: 15 percent
This coefficient is tested for just a single value - 5 percent, which is materially lower than the base model value. The discount assumption is obviously important but does not have a dominating influence on model outcomes.

Table 18 Foreclosure discount impact on risk weights

Bucket \ Discount	*15 percent*	*5 percent*
40/90	*143*	*120*
30/70	*28*	*23*
20/60	*6*	*5*
Portfolio	*24.7*	*20.9*

5.4 'Negative' and 'positive' scenarios

As well as the single coefficient sensitivity tests set out above the model was also tested with positive and negative scenarios where a range of input coefficients were systematically skewed from the base run numbers.

In the negative scenario the nine variables which were tested above were altered by ten percent in a direction which adds to measured risk. In the positive scenario the variables are altered by ten percent in a direction that reduces risk. Outputs for selected risk buckets for both scenarios are presented in table 19. While a ten percent change in a particular input may seem small given the uncertainty around a particular coefficient value, the requirement that every coefficient changes by ten percent is quite a stern test. It would be expected that there would be unders and overs from our best estimates of the true value of the coefficients and the likelihood that every one of the nine coefficients was out by the same percentage amount would be relatively small. While it is not possible to give the difference between the positive and negative scenarios a strict quantitative interpretation, it can be viewed as a rough confidence band around the point estimates represented by our base scenario. These results suggest that this 'confidence' band is reasonably wide and is relatively larger for the lower risk buckets.

Table 19
Positive and Negative Scenario impact on capital

Scenario / *Bucket*	*Benchmark*	*Negative scenario*	*Positive Scenario*
40/90	*143*	*225*	*86*
30/70	*28*	*48*	*14*
20/60	*6*	*12*	*3*
Portfolio	*24.7*	*40.0*	*15.0*

5.5 Sensitivity test conclusions

The sensitivity tests lead us to the following conclusions about the robustness of the model.

1. The model inputs can be separated into three classes according to their impact on the models risk weight outputs.

   Low impact: Demographic coefficient, foreclosure discount

   Medium impact: Unemployment volatility, sensitivity to debt servicing.

   High impact: House price volatility; interest rate volatility; house price/interest rate correlation.

2. Of the high impact variables the size of interest rate/house price correlation coefficient is critical. A strongly positive correlation takes much of the risk out of floating rate residential housing portfolios.

3. The model does not have knife edge properties either with respect to a particular input or a cluster of inputs.

4. A reasonably wide range of risk weight output can nevertheless be generated by moderate changes in key inputs.

6 Comparison with US Federal Reserve and APRA modelling

This section compares the main TUI outputs with relevant outputs presented in the Federal Reserve (Calem and Follain's (2003)) analysis of the Basel II IRB residential lending model and the stress test of residential lending portfolios generated by APRA's project Panama.

6.1 US Federal Reserve Model Analysis

Two sets of results from the Calem and Follain study are presented below. The first, the risk weights for a nationally diversified portfolio, is shown in table 20[1]. The second, which represents a regionally concentrated portfolio, is shown in table 21. The buckets portfolios are differentiated by their LVRs and by their FICO scores. A FICO score is a proprietary creditworthiness score with higher numbers representing a better credit.

The regional results may be more relevant as a New Zealand comparator because the diversification possibilities in a smaller geographical area may be closer to that of the New Zealand housing market. The national results tend to reflect the limitations of the Federal Reserve modelling process which is grounded in empirical outcomes. Like TUI, the model builds up a loss distribution but it does so by drawing on the actual US historical experience. Over the period the model was estimated, it turned out there was no instance in which there was a material fall in the national house price index. The low risk weights in the 'quality' part of the table reflect that experience. On the other hand there were instances of significant falls in a few regional housing markets and the regional portfolio risk weights reflect those cases. We would expect that if the model were rerun in a couple of years then the national portfolio risk weights would be much higher than the current estimates reflecting the sharp downturn in prices the US is now experiencing.

Table 20
US National portfolio risk weights

LVR / ***FICO***	***95***	***90***	***80***	***70***
620	*45*	*40*	*33*	*28*
660	*39*	*15*	*24*	*20*
700	*20*	*13*	*8*	*8*
740	*8*	*5*	*3*	*3*

1 Note that the US results have been calibrated differently to Basel II outputs in an important respect. It is assumed, logically, that only tier one capital (a risk weighted minimum capital level of four percent) is available to prevent default. The Basel IRB model, on the other hand, assumes that all of the eight percent minimum capital is available to absorb credit losses. The TUI results have also been presented on the Basel II basis. To compare the US results with the TUI results it is necessary, therefore, to divide the US results by two. The 700 FICO /90 LVR bucket risk weight becomes 13 rather than 26. All of the results in tables 20 and 21 have been presented on this adjusted basis.

Table 21
US regional portfolio risk weights

FICO \ LVR	95	90	80	70
620	*120*	*89*	*68*	*53*
660	*87*	*63*	*46*	*35*
700	*38*	*25*	*19*	*13*
740	*13*	*8*	*5*	*4*

To compare the US and TUI results we need to make some assumption about the comparability of FICO and DSR scores and we have assumed that the higher FICO score risk buckets (at the bottom of the table) have some rough comparability, from a risk perspective, to TUI's moderate (say around 20-30 percent) DSR tranches.

The most relevant comparison is between TUI's fixed rate mortgages outputs (see section 7) presented in table 22 and the US regional data.

The comparison suggests that the TUI results are not out of line with the US results. And as the latter are the lowest of the four sets of results, from a range of models presented in the Calem and Follain paper, this further suggests that TUI's unemployment and demographic risk drivers are not overstated and/or that macro assumptions used in the TUI base run are not excessive.

Table 22
New Zealand long term fixed rate mortgages

DSR \ LVR	100	90	80	70	60
40	*29*	*28*	*23*	*14*	*5*
30	*29*	*28*	*23*	*14*	*5*
20	*20*	*18*	*15*	*9*	*3*
15	*10*	*9*	*8*	*5*	*2*

Both sets of results show similar sensitivities to the original loan to valuation ratio. The difference between the 70 percent LVR and the 95 percent is nearly 3:1 for the US for the 700 FICO buckets. The comparable ratio for New Zealand was a little over 2.

One of the differences between the FED model and TUI results is the relationship between stressed LGDs and LVRs. The US National stressed LGDs range from 49 percent for the 95 percent LVR tranche to 21 percent for the 70 percent LVR tranche. The TUI stressed LVRs, on the other hand, range from about 45 percent to 35 percent over the 95-70 LVR buckets.

The US Fed stressed LGD are also, on average, lower than the TUI stressed LGDs. In part the difference is explained by the following:

- TUI uses a higher discount rate than the Federal Reserve Model;

- The Federal Reserve stressed LGDs are calculated from the most stressed period in their historic simulation period but because this did not include a period when there was a pronounced fall in house prices the foreclosure discounts embedded in the data may not have been high as assumed for the TUI modelling.

6.2 APRA's project Panama

It is not possible to make a direct comparison between the TUI and APRA project Panama capital because the Panama exercise does not directly generate capital ratios. Rather Panama generates losses for a range of loan instrument risk characteristics, mainly distinguished by their loan to valuation ratios, from a single stress event. This event is a thirty percent fall in house prices which generates an overall default rate of three percent and an average loss rate of around one percent.

While the Panama stress event does not detail the increase in the interest rate and unemployment rate that would have to accompany the decline in house prices to generate the estimate for the level of losses it appears that the implied Australian stress event is similar to the UK event that has been used to calibrate TUI. We can then reasonably compare the Panama PDs and stressed LGDs with the stressed PDs and LGDs generated by TUI.

Panama generates a strong relationship between LVR and the probability of default. The difference between the 96-100 percent LVR bucket and the 71-75 LVR bucket was around 130 percent. This compares with the TUI base case margin of 225 percent.

Panama's LGD profile is different to TUI's. It ranges from nearly 50 percent with a 100 percent LVR to only 5 percent at a 60 percent LVR. TUI generates a flat relationship.

Panama's average LGD is 28 percent compared to 42 percent for TUI. About half of this difference is because Panama does not take account of the time value of money. The Panama empirical work showed little relation between LVR and downturn LGD but this was rejected in favour of a set of 'adjustment factors' which embedded a strong relationship between LGDs and LVRs which were provided by a mortgage insurer to the Basel Committee in 2001

Taking the PD and LGD results together, and adjusting for the time value of money, it appears that TUI and Panama will generate similar losses and capital requirements for the higher LVR buckets (over 75 percent).

In the lower LVR buckets of 75 percent and less, which account for a significant portion of most banks loans, there are substantial difference between TUI and Panama which increases as the LVR decreases, For example even after adjusting for time value the LGD for Panama's under 60 percent LVR is 12 percent compared to TUI's 43 percent.

As noted above we are reasonably confident about the TUI estimates because they necessarily follow from a transparent set of assumptions. On the other hand the figures which underpin

the APRA estimates were obtained from a third party source and we have been unable to ascertain why they differ so materially from the TUI results.

The differences between the Panama and TUI LGD outputs are material. If it is assumed that the TUI figures are correct then the average loss rate from the Panama stress test (even

excluding the time value effect) would have been more like 1.5 percent of loan assets rather than the reported 1 percent.

7 The Full TUI Model

This section describes a set of modules that can be added to the basic TUI model. The modules are:

1. A **'rational forbearance'** module that deals with alternative ways of dealing with mortgagee default.

2. An **interest rate duration** module that allows the analysis of fixed interest rate portfolios with varying durations.

3. A **mortgage insurance** module that analyses the impact a range of mortgage insurance arrangements and different assumptions as to the effectiveness of the insurance on capital requirements.

4. A **self employed** module that replaces the employed borrowers in the core model with self employed borrowers.

5. An **investment module** that deals with portfolios of loans for rental housing. This module is undeveloped at present.

Figure 5 Structure of the full TUI model

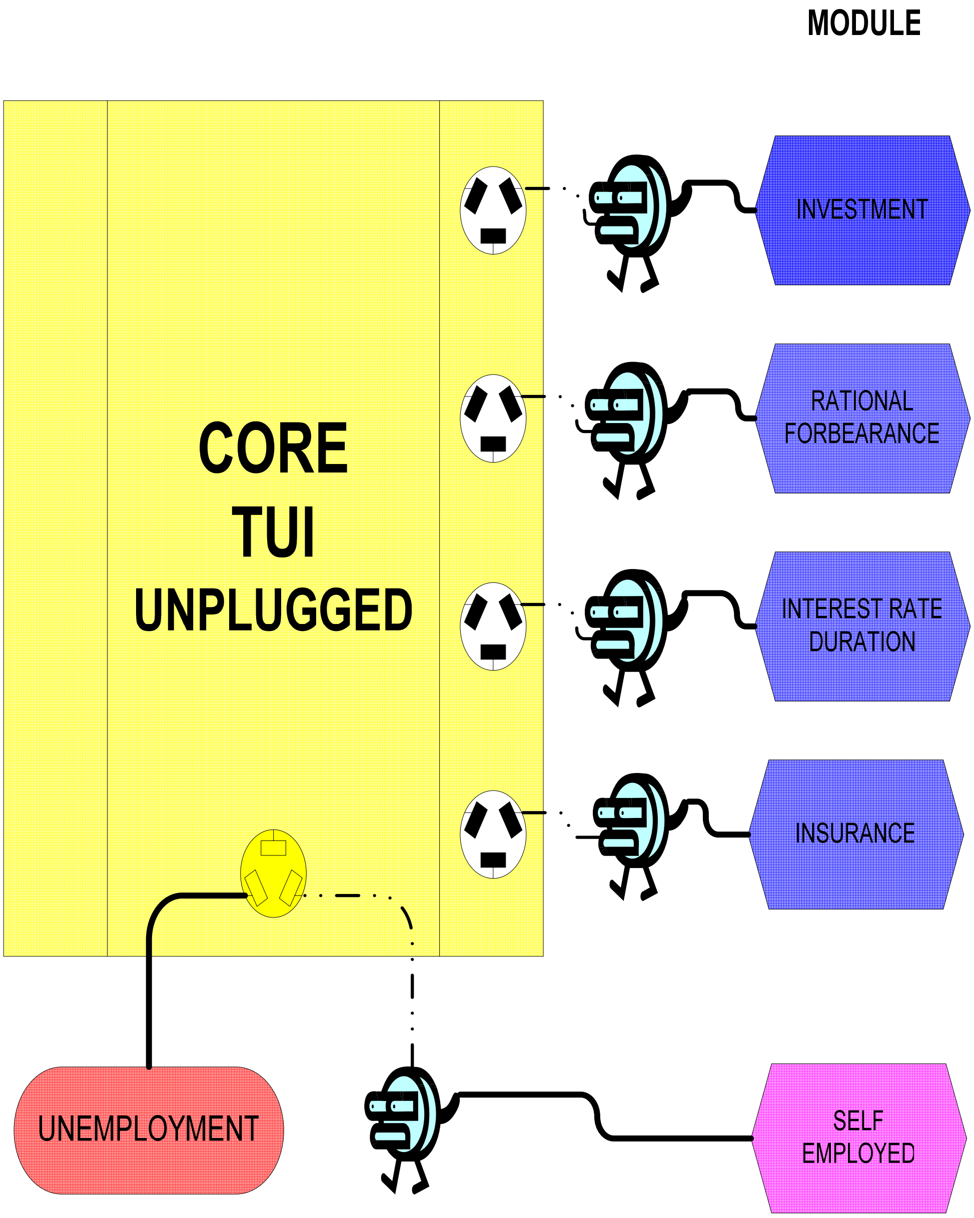

7.1 Insurance Module

The insurance module assesses the impact of different insurance arrangements on the capital requirement. In particular the different typically 'Australian' and typically 'New Zealand' insurance arrangements are captured and allowance can be made for the failure of the insurer to make a full payout on claims.

In Australia the typical insurance arrangement covers the full value of the loan. The only component of economic loss that is not covered is the time value of money. In New Zealand the coverage is typically less comprehensive. Only the funds advanced over 80 percent of the value of the property are insured rather than the whole amount of the loan. These different contracts are referred to here as 'Australian' and 'New Zealand' contracts.

The inputs into the module are:

- Whether the risk bucket is insured.
- Whether the contract is an ‘Australian’ or a ‘New Zealand’ contract.
- The probability that the insurer defaults when a payout is demanded in a stress situation.
- The proportion of claims met by the insurer when it does not default.

The later two inputs merit some discussion.

Probability insurer fails

The relevant input from a capital adequacy perspective is the probability that the insurer will default when the bank is in a stressed situation and is making large claims. It is the tail events that drive the capital requirement. Because any individual bank stress event is likely to be, at least, an Australasian wide event, and probably a world wide event, then the likelihood of the joint event of a large claim and insurer failure will be much higher than if the two events were independent. The joint probability is also likely to be higher again when the insurer is an Australian incorporated insurer because the diversification effect will be less than with a large, internationally diversified entity.

The probability of failure has been set judgementally (the mid point estimate is ten percent) in the illustrative runs but a further development of the model could be a structural approach to the probability of failure of the insurer. The insurance company balance sheet could be an explicit variable and the claims on the capital would be the counterpart of claims by banks in a stress situation.

If the insurer fails it is assumed that the payout is zero. While there could eventually be some positive payout, at the point of the failure auditors and directors are likely to take a very conservative view of the value of the claim for the purposes of assessing the solvency of the bank. We have assumed here that its value would be completely discounted.

Insurer failure is modelled as follows: For each run of the macro engine TUI draws a random variable (independent of all macro-variables) from the distribution $U\,[0, 1]$. If this variable is less than the probability of default then the insurer defaults and the insurance payout is set to zero. If the insurer does not default then the insurance module goes on to calculate how much the insurer pays.

Proportion of claims paid

In a bank stress situation it is likely that the proportion of claims paid will fall. The insurer is likely to pay greater attention to whether the conditions of the loan contracts have been strictly complied with. Claims that might have been met in ordinary times may be rejected in a stress situation and the insurer may even ‘invent' problems to slow their cash payouts. The proportion of claims which are declined in a stress situation has been set, judgementally, at 20 percent in the base case.

The total expected payment shortfall of 30 percent is consistent with a rule of thumb used by one New Zealand bank.

7.2 'Rational Forbearance' Module

One reaction to the losses generated by the TUI LGD module is that it exaggerates likely losses. The reason is that in an acute stress situation banks would not automatically foreclose on their security, evict the mortgagee and sell the property. Rather, in many cases they would leave the mortgagee in possession providing he was making a reasonable effort to service the loan. Because of the dead weight costs of foreclosure there are opportunities for both parties to gain by such an arrangement. These gains include:

- Real estate and other direct costs of a mortgagee sale.
- The impact of a repossession mortgagee sale on the bank's reputation.
- The impact of a mortgagee sale on the realisation at sale. TUI assumes a discount of fifteen percent.

The rational forbearance module allows the bank to exploit some of these possible gains by allowing the borrower to remain in possession of the property where there is a net advantage to the bank in doing so.

The adjustment for rational forbearance is calculated by multiplying the average loss given default by the proportion of defaults to be handled by forbearance rather than by foreclosure, and by the proportion of losses saved by the forbearance process. This amount is netted from the recorded loss.

This module is normally switched off. We have doubts as to whether rational forbearance, on any significant scale, would actually work. Many bankers think that forbearance strategies could end up costing more than early action. Secondly, we have little idea as to what the relevant coefficients might be. The module is there for users who might have a better feel for the relevant inputs and who would like to explore what impact different assumptions might have on capital requirements or on stress event scenarios.

7.3 Self employed Module

The self employment module calculates the capital requirements of loans to self employed borrowers. The difference from the employed borrower in the base model is that the unemployment variable is replaced by volatility in self employed incomes. The self employment volatility figure has two components; the volatility in aggregate self employed income and the conditional volatility of individual incomes. The first is assumed, in the New Zealand context, to be equivalent to the volatility of the operating surplus component of GDP. The second component is estimated from an analysis of longitudinal data of individual incomes.

This module is in a preliminary build stage pending the availability and analysis of relevant and reliable data.

7.4 Rental housing investor module

This module is also in a preliminary build stage. There are a number of problems in developing a simple model.

- There is a multiplicity of investor circumstances or business models. These can range from the investor with a good salary income and a high-value, mortgage-free house which is providing security for a modest rental investment to a highly leveraged position by a professional investor with no, or limited, financial resources beyond the cash flows from those properties.
- Banks may not know which of their exposures are investments. Some may have started out as owner occupied but may have changed their status.
- We think that investor behaviours could be different from that of owner occupiers but we do not know how.
- Investment dwelling prices could behave differently from owner-occupied dwelling prices. It is difficult to get data on this behaviour because house price data does not distinguish between the two types of housing.

One way forward could be to create a number of stylised investor types and to model them separately. This would help with the heterogeneity problem noted above.

Despite these difficulties it is possible with the current model to make coefficient adjustments that might capture some of the risk characteristics of resident investment lending. The first is to increase the volatility of either, or both, of the housing index and individual house price distributions. The logic here is that there are likely to be sharper decreases in the value of rental housing in a stress situation when investors are forced to unload their investments.

A second possibility is that the sensitivity of the rate of stressed sales to increases in interest rates could be set to be greater for investors than for owner occupiers. This would reflect the greater commitment of owner occupiers to hang on to the family home compared to the investor’s commitment to his investment portfolio.

7.5 Fixed Interest Module

It is easy to model a long term mortgage instrument similar to the US 30 year fixed rate instrument. The interest rate volatility coefficient is set to zero.

However, analysing the risk of the shorter tem fixed rate loans (six months to three years) that are marketed in New Zealand is more difficult. This involves assessing the effects of a

partial, rather than a full period, impact of an increase in debt servicing costs. TUI was not constructed to answer questions about what happens within its observation period.

However, to give some sense about the difference New Zealand fixed rate mortgages might make to capital requirements the following, fairly rudimentary, fixed rate module has been devised.

The key assumptions underpinning the module are as follows.

- The fixed rate portfolio is evenly spread over time. A two year portfolio, for example, will have half its loans with less than one year to maturity and half with one to two years to maturity.
- The interest rate reset point is the middle of the year.
- The impact of an interest rate shock on stressed sales is distributed in the following manner over the three year observation period. 20 per cent is in the first year; 60 percent in the second, and 20 percent in the third year. The logic here is that relatively few borrowers default in the first year because they draw on whatever resources they have to avoid a default. The major impact happens in the second year as these resources become exhausted. The third year picks up the residual. It is possible to change these proportions within the fixed interest module.
- The impact on the probability of stressed sales is proportionate to the time the interest rate shock is endured and the above year by year weightings. Thus the effect of a fixed rate mortgage can be calculated by adjusting the linear stressed sale sensitivity coefficient.
- On the above assumptions the DSR coefficients adjustments are as follows:

Table 23 Fixed interest adjustment factors

Portfolio	***Percentage of base***
One year	*90*
Two years	*70*
Three years	*50*

7.6 Fixed Interest portfolio results

Table 24 shows the risk weight for selected risk buckets for the base case; the one, two and three year fixed rate mortgage portfolios; and a long term fixed rate mortgage portfolio. The long term results assume that there is no interest rate impact at all on the existing pool of mortgages.

It is clear that long term fixed rate mortgages make a substantial difference to the risk of a residential mortgage portfolio. The obvious reason is that interest rate volatility is the key driver of stressed sales risk and that removing this driver will materially reduce risk. Note that this reduction in risk is greatest where the borrower's debt servicing ratio is already

burdensome and reflects the way the model has been calibrated to reflect the greater vulnerability of this group to interest rate shocks.

With the one, two and three year fixed rate portfolio there is a moderate reduction in risk with the higher debt servicing ratios but this reduction drops away as the debt servicing ratio falls. Given the relatively short average duration of New Zealand mortgage portfolios the impact on the risk of banks' overall portfolios of the interest rate fixing that does occur appears to be moderate.

Table 24
Fixed rate mortgages summary

	Floating	*One year*	*Two year*	*Three year*	*Long term*
(40/90)	143	130	105	80	28
(30/70)	28	27	24	21	14
(20/60/)	6	6	6	6	3
Portfolio	24.7	23.1	19.8	16.6	8.2

8 Basel II IRB applications

In this section we test whether the correlation coefficient of 0.15 in the IRB housing equation is appropriate to New Zealand circumstances and set out a model which translates TUI's PD and LGD outputs into outputs which are consistent with the Basel II IRB definition of default.

8.1 Testing the Basel II IRB housing equation correlation coefficient

The test of the IRB correlation coefficient involved putting TUI generated inputs for long run PD and stressed LGD into the IRB housing equations to generate an array of risk weights by LVR and DSR. The results are set out in table 25. For ease of reference the comparable base run data is represented in table 26.

It is clear from tables 25 and 26 that TUI generates materially higher risk weights than the Basel II IRB model for the same PD and stressed LGD inputs. This means that the implied TUI correlation coefficients must be higher than the IRB mandated number of 0.15. These implied correlations, which are calculated by adjusting the correlation coefficient in the IRB equation until the risk weight equals the base case risk weight, are set out in table 27.

The key question from a capital perspective is which of these correlation estimates is right for New Zealand - the IRB 15 percent figure - or the array of TUI estimates.

We understand that the IRB figure was originally set by 'reverse engineering' a sample of banks' economic capital numbers to derive an implied correlation coefficient. The US

Federal Reserve study (Calem 2003) provides some support for this figure though it also cites some results that suggest a slightly higher number might be justified.

Table 25
Basel II IRB model with TUI Inputs

LVR / DSR	*100*	*90*	*80*	*70*	*60*	*50*
40	*68*	*52*	*34*	*21*	*11*	*9*
30	*53*	*40*	*25*	*15*	*7*	*5*
20	*37*	*29*	*18*	*10*	*5*	*5*

Table 26
Risk weights by risk bucket base case

LVR / DSR	*100*	*90*	*80*	*70*	*60*	*50*
40	*188*	*143*	*88*	*46*	*18*	*6*
30	*101*	*79*	*52*	*28*	*11*	*4*
20	*52*	*45*	*33*	*18*	*6*	*2*

From a New Zealand perspective the significance of the US results is that if they are correct (and there is emerging evidence that they are too low) then they apply to long term fixed rate mortgages only. The balance between systematic and idiosyncratic risk which the correlation coefficient captures can be significantly different for floating rate and short term fixed rate mortgages. Floating rate mortgages are subject to the additional systematic risk that interest rates will increase. We would expect, therefore, the correlation coefficient for floating rate portfolios to be higher than the coefficient for fixed rate portfolios.

Table 27 sets out the implied TUI coefficients for an array of risk buckets. The two main features of this analysis are first: the average coefficient of about 0.25 is substantially higher than the Basel figure: and second the correlation coefficient is a function of risk. This accords with intuition. For high LVR buckets the borrower is very exposed to movements in the average house price index. With lower LVR tranches on the other hand, a combination of systematic and an idiosyncratic house price shock is often necessary to induce a default.

Table 27
TUI Implied Correlations

LVR / DSR	**100**	**90**	**80**	**70**	**60**	**50**
40	*0.38*	*0.37*	*0.33*	*0.29*	*0.23*	*0.17*
30	*0.27*	*0.27*	*0.28*	*0.25*	*0.21*	*0.15*
20	*0.20*	*0.22*	*0.20*	*0.22*	*0.18*	*0.15*

8.2 Converting TUI outputs to Basel II outputs – Introducing Basel-TUI

A second exercise is to use TUI to assess the reasonableness of banks' estimates of the key inputs into the Basel II housing equation. This is not a straightforward exercise, because, as **noted** in section 2 TUI uses an economic definition of default whereas the Basel II definition is based on an accounting concept. This means that measured PDs and LGDs are different and that the TUI measures of long run average PD or stressed LGD cannot be directly compared with and substituted for bank estimates of these variables.

To generate IRB consistent PDs and stressed LGDs a second model called Basel-TUI has been developed. Basel-TUI does not change any of the economic or behavioural processes in the underlying TUI model. It simply transforms the TUI outputs, which are based on an economic concept of default (all losses are insolvency events), to be consistent with the Basel default definition (the 90 day overdue rule captures both solvency and liquidity events).

In this section we first describe the logic underpinning Basel-TUI. We then describe some key properties of the model's outputs and produce some summary stressed LGD and risk weight results for an illustrative portfolio. The final part draws some key conclusions.

The Basel-TUI logic

The modelling starts with the relationship between the Basel PD definition and the stressed sale (SS) concept, which is a critical driver of the TUI results. The two differences between the measures (see Figure 6) are first: the SS concept includes loans where the home is sold before the loan goes into default (without a loss to the bank), which the Basel PD excludes; and second, only Basel PD includes loans that go into default but which are subsequently fully cured.

Broadly, the factors which drive stressed sales will also drive Basel PD. An increase in unemployment, for example, will mean that some borrowers will not be able to service the loan and will either sell the house or the bank will foreclose. These borrowers are captured by the stressed sale concept. In the same shock other borrowers will not be able to service their loan for more than 90 days but will find another job and the default will turn to a cure. These are captured by the Basel PD definition.

As Basel PD can be described as a function of SS it follows that if we can model this relationship then we can generate the stressed LGD and long run PD inputs required by the IRB risk weight equation. However, it turns out that this is not a simple exercise because the relationship between IRB PD and SS is a complex non-linear one and we have little empirical evidence of its structure. We have not, therefore, been able to develop a convincing and robust model which generates stressed LGD estimates. What we have done instead is to focus more on generating a framework that will allow supervisors and banks to input the key assumptions to generate Basel IRB consistent LGDs, PDs and capital numbers in a consistent and transparent manner.

The Basel-TUI framework

The Basel-TUI PD/LGD estimation framework is based around the default categories set out in Figure 6 and requires just two inputs.

The first is ALPHA. This is the ratio of sales of the house by the customer (without loss to the bank) after 90 days in arrears to total sales of the house by the customer – i.e. the ratio of category (2) to categories (1) and (2). Our simulations show that the overall portfolio capital ratio is not very sensitive to changes in this ratio. For example, a change in the stress event ratio of 50 percent to 40 percent would change the risk weights by only 0.2 percentage points on average. Because of this lack of sensitivity we have assumed a fixed ratio of 50 percent for the rest of the analysis so we can focus on the second, more important, ratio, GAMMA.

Figure six
Relationship between stressed sales economic default and Basel IRB default

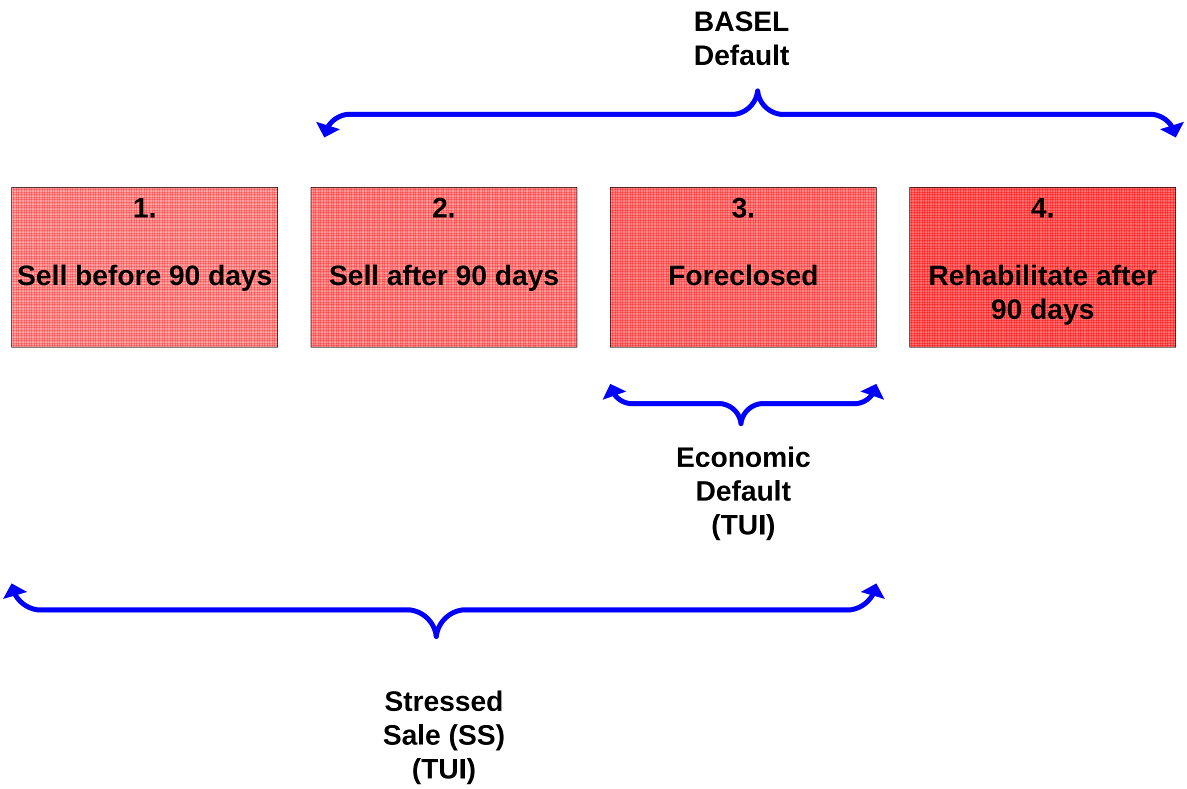

GAMMA is defined as the ratio of the Basel defaults in category four to those in categories two, three and four. GAMMA is can be thought of as a cure rate, though it might be somewhat different to what some bankers refer to as a cure rate.[2]

[2] Note that it appears that bankers sometimes use the term cure rate in a Basel II context to include all Basel II defaults that do not involve a loss of principal. They may therefore include both categories 2 and 4 in figure 6.

The key relationship in the framework is that between TUI stressed LGD (TSLGD) and the Basel II stressed LGD (BSLGD).

BSLGD = TSLGD * TD/ BD

Where	TD is the TUI default rate in the stressed scenario.
and	BD is the Basel default rate in the stressed scenario

The higher the ratio between the TUI default rate and the Basel default rate the higher will be the observed LGD in the stress situation. The ALPHA and GAMMA inputs work by impacting on this ratio.

For example, assume that the TUI stressed LGD is 50 percent; that the TUI default rate is 2 percent; and that (for simplicity) no loans fall into categories one and two (this would be close to the situation with high LVR loans with a 30 percent house price decline). If GAMMA were 30 percent then the observed Basel default rate would be 2.9 percent and the Basel-TUI stressed LGD would be 50 * 2/2.9 = 34.5 percent.

Note that for simplicity we have assumed that the liquidity events represented by categories two and four do not generate any losses. That is equivalent to assuming that banks' penalty rates are sufficient to cover the time value and administrative costs associated with a pure liquidity event. The banks, however appear to be assuming that there will be positive costs. The model outputs can be adjusted to account for these costs but the effect on the figures reported below would not be large. The adjustment would depend on banks' assumptions about costs, discount rates and penalties and is likely to add between 1 and 2 percentage points to the Basel-TUI stressed LGD figures.

The Estimation Process

The estimation process has the following steps.

- TUI generates a base run which provides a set of long run average economic PDs which are differentiated with respect to LVR.

- A stress scenario is selected and TUI generates economic default rates (TD) and loss rates (TSLGD) are generated for the scenario.

- Assumptions are made as to the relationship between the components of TUI SS rate and the Basel default rate in an unstressed environment. This gives a starting point default rate which inputs into the calculation of the average default rate. We have drawn on banks' empirical work here.

- Input ALPHA and GAMMA into Basel-TUI and generate Basel long run PD, stressed LGD and associated capital ratios.

Selection of the stressed scenario

The benchmark scenario, which generates a loss that is consistent with the 0.999 solvency standard in the Basel model, is as follows:

House price index - 30 percent fall
Interest rates - 3.75 percentage point increase
Unemployment - 4.5 percentage point increase

The key variable in the scenario is the house price index fall assumption. It directly drives almost all of the changes in stressed LGD whereas the role of the other two variables is largely to provide the context in which the reasonableness of the GAMMA input can be assessed.

Besides its consistency with the Basel solvency standard, the 30 percent house price fall was chosen because:

- It is consistent with the Australian FSAP stress test house price fall assumption.
- It is consistent with the Project PANAMA house price fall assumption.

Model outputs

The following set of results illustrates the key properties of the model.

(i) Relationship between stressed LGD, GAMMA and LVR

The Table 28 results show, first, that stressed LGD decreases as GAMMA increases. This result follows directly from the stressed LGD equation presented above. The higher the GAMMA, the higher the proportion of liquidity events (with no losses) in the total level of defaults and hence the lower the stressed LGDs. Second, stressed LGD decreases as LVR decreases. This reflects the higher proportion of economic defaults in the stressed situation with high LVR loans. With lower LVR loans many, and sometimes most, owners will still have positive equity after the house price shock and will be able to repay their loan when they encounter servicing difficulties.

Table 28
Stressed LGDs for different LVRs and GAMMAs

LVR / ***GAMMA***	***100***	***90***	***80***	***70***	***60***	***50***
10	*48.3*	*43.8*	*37.8*	*27.6*	*12.6*	*5.1*
20	*42.9*	*39.0*	*33.6*	*24.5*	*11.2*	*4.3*
30	*37.5*	*34.1*	*29.4*	*21.5*	*9.8*	*3.8*
40	*32.2*	*29.2*	*25.2*	*18.4*	*8.4*	*3.2*
50	*26.8*	*24.4*	*21.0*	*15.3*	*7.0*	*2.7*

(ii) Relationship between GAMMA and long run PD

The positive relationship between GAMMA and the long run PD shown in table 29 reflects the assumption, embedded in a higher GAMMA, that there are a higher number of loans which experience temporary servicing problems in the stress situation. This, in turn, increases the measure of loans experiencing temporary liquidity problems in all other scenarios thus increasing the long-run average PD. Figure 7 illustrates the relationship showing the impact of an increase in GAMMA on the stress event default rate and on the long run average default rate.

Table 29
GAMMA and long run PD

GAMMA	*Long run PD %*
10	*1.25*
20	*1.33*
30	*1.43*
40	*1.56*
50	*1.74*

The most plausible GAMMA is probably in the 30-50 percent range suggesting a long run PD of around 1.5 percent. This is consistent with the Basel II quantitative impact study (QIS5) which reported an average PD of 1.5 for class 2 banks.

Figure seven
Impact of increase in GAMMA on PD

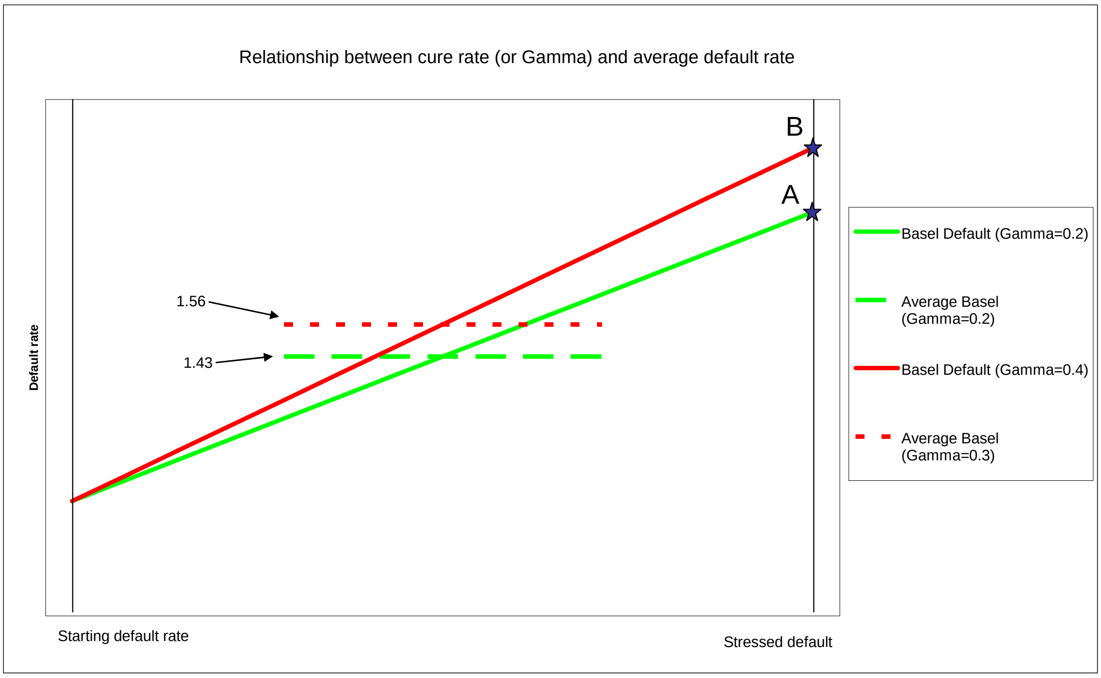

(iii) Relationship between long run PD, DSR and LVRs

Table 30 shows that there is only a relatively weak relationship between LVR and Basel PD. This contrasts with the economic model where there is a much stronger relationship. This result is also broadly consistent with New Zealand banks' analyses, which show little or no relationship between LVRs and Basel default rates. The relationship between risk and LVR is still there in the Basel model but much of the empirical representation of it has migrated to the stressed LGD measure.

The relationship between default rates and debt servicing ratios, on the other hand, is more pronounced. However, this may overstate the real-world relationship because servicing capacity is a more complex variable than that captured by the simple TUI DSR variable.

Table 30
Long run default rates % (GAMMA 0.3)

LVR / DSR	***100***	***90***	***80***	***70***	***60***	***50***
40	*2.9*	*2.9*	*2.9*	*2.6*	*2.0*	*1.8*
30	*1.9*	*1.9*	*1.9*	*1.7*	*1.4*	*1.3*
20	*1.0*	*1.0*	*1.0*	*0.9*	*0.9*	*0.8*

(iv) Impact of GAMMA and LVR on risk weights

The array of risk weights in the top section of table 31 were generated by inputting the PDs and stressed LGDs calculated by Basel-TUI into the IRB housing risk weight equation.

For comparative purposes two other sets of risk weights generated by TUI are also shown. The first are the weights generated when TUI PD and LGD estimates are inputted into the IRB risk weight equation. The second set are the risk weights generated directly by the TUI model. The latter are higher those generated by the IRB equation because, as noted in 8.1 above, the Basel correlation of 0.15 appears to be too low for New Zealand portfolios.

Table 31
Risk weights for different LVR and GAMMA combinations
(Assumes 30 percent DSR)

LVR / GAMMA	***100***	***90***	***80***	***70***	***60***	***50***
10	*82*	*75*	*64*	*44*	*18*	*7*
20	*77*	*70*	*60*	*41*	*17*	*6*
30	*71*	*65*	*56*	*38*	*15*	*6*
40	*65*	*59*	*51*	*35*	*14*	*5*
50	*59*	*54*	*46*	*31*	*12*	*5*
60	*52*	*48*	*41*	*28*	*11*	*4*
TUI inputs to Basel Model	*53*	*40*	*25*	*15*	*7*	*5*
True Model TUI	*101*	*79*	*52*	*28*	*11*	*4*

The table shows that risk weights decrease as GAMMA increases. However, risk weights are not strongly sensitive to GAMMA because of the offsetting impacts of the changes in stressed

LGD and long run PD estimates on the risk weight. A higher GAMMA reduces stressed LGDs, which decreases the risk weight, but also increases long run PD which partially offsets that decrease.

The table and figure 8 also show that the relationship between the true (TUI) risk weight and the Basel II risk weight depends on LVR. With a GAMMA of 0.3 TUI generates higher risk weights with LVRs which are over 80 percent. With lower LVRs the Basel risk weights are higher. The explanation for this relationship is that there are two distortions at work in the Basel risk weight equation. On the one hand the inclusion of liquidity events in the default definition overstates risk. On the other, the correlation coefficient is understated and the extent of the understatement is highest with the high LVR tranches. For low LVRs the former distortion is dominant. For high LVRs the correlation effect outweighs the distortion caused by the PD definition.

Figure eight
Relationship between risk weights, GAMMA and LVR

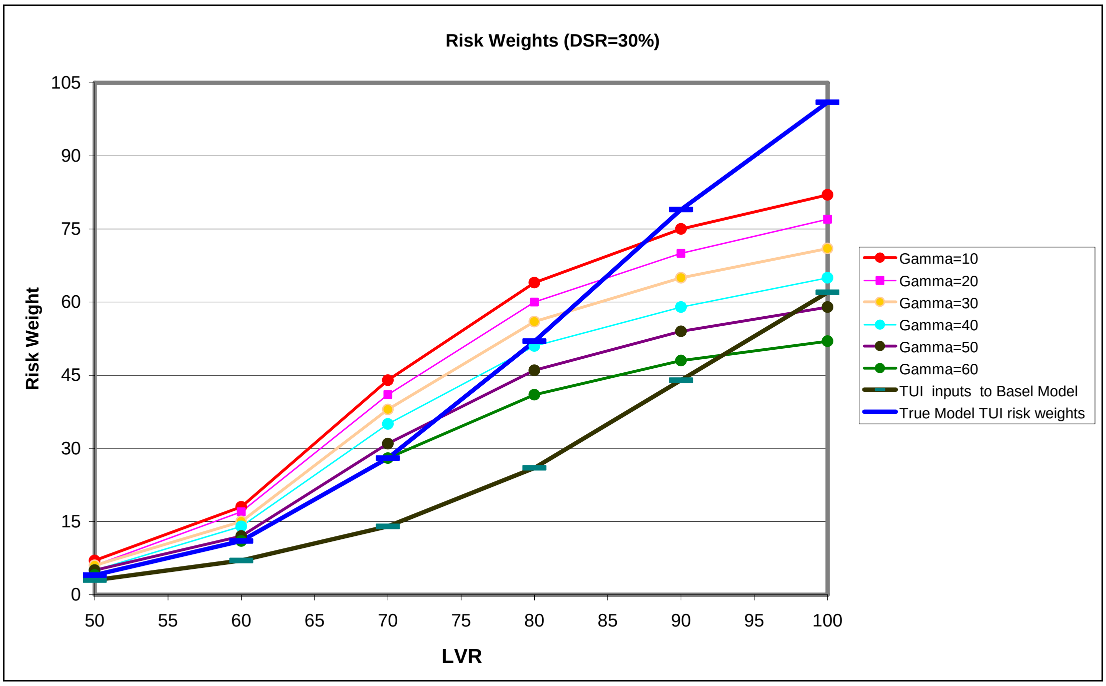

(v) Portfolio average stressed LGDs and risk weights

Figures 9 and 10 show, for our illustrative portfolio (based on 2004 HES data), the impact of different GAMMA assumptions on the portfolio average risk weight and stressed LGD.

The figures capture the combination of assumptions concerning the use of insurance and the application of the 10 percent LGD floor. With respect to the insurance it is assumed that that loans with an LVR of 80 percent and above are insured and that the stressed payout rate is 80 percent.

Figure nine
Relationship between GAMMA and average stressed LGDs

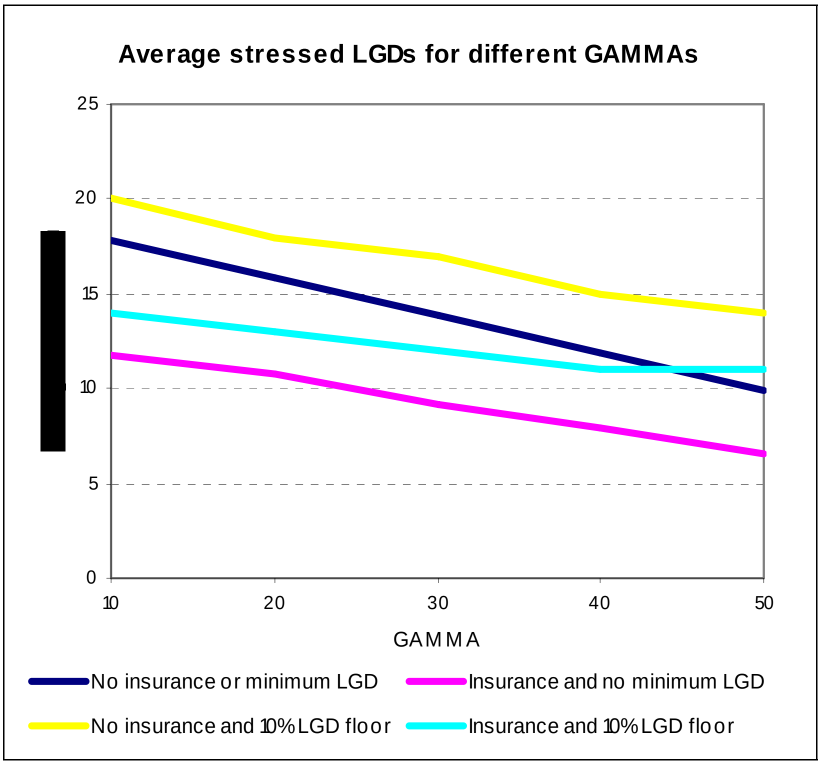

Figure ten
Relationship between Portfolio risk weights and GAMMA

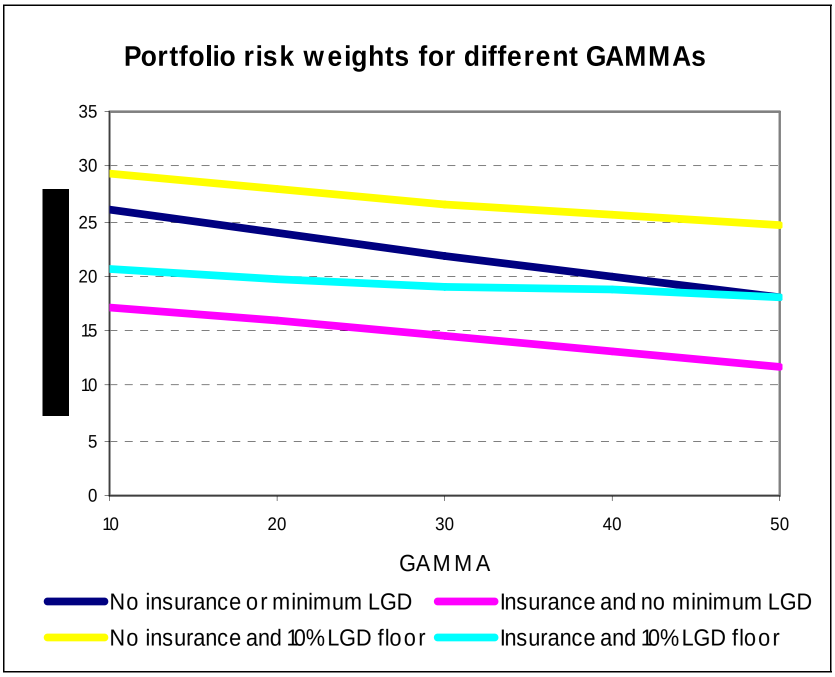

Table 32 provides a summary of the results assuming a thirty percent GAMMA.

Table 32
Results summary assuming 30 percent GAMMA

	Portfolio Risk weight	*Average Stressed LGD*
No insurance No minimum LGD	*22*	*14*
Insurance No minimum LGD	*15*	*9*
Insurance Minimum LGD	*19*	*12*
No insurance Minimum LGD	*27*	*17*

The LGD floor increases the average stressed LGD by around 3 percentage points because the 10 percent minimum replaces the lower estimated LGD for low LVR risk buckets. Note that these figures do not include an estimate of LGDs on cured loans and are therefore understated by around 1-2 percentage points.

A summary of key results

- LVR is a strong risk weight driver.
- Risk weight results will depend on the GAMMA assumption but are not unduly sensitive to it.
- The 10 percent LGD floor increases the portfolio average stressed LGD by around 3 percentage points.
- The Basel default definition introduces a distortion that leads to risk weights being overstated. This provides an offset to the understatement due to a low correlation coefficient for low LVR tranches but does not entirely compensate for the correlation problem for high LVR tranches.

9 TUI in simulation mode

This section presents the results from a set of simulation exercises. The simulations are conducted by setting the macro volatility figures at zero and entering the macro values that represent a particular scenario as the expected outcome values. TUI's behavioural structure and idiosyncratic house price generator will then generate the default and loss outcomes for that scenario. The scenarios presented here are:

1. A replication of 'normal' and 'good time' housing default scenarios for New Zealand.
2. A replication of relevant New Zealand FSAP stress tests.
3. A replication of the New Zealand mid-1980s experience. This is a case where there was a substantial increase in interest rates which did not generate large housing losses because it was accompanied by stable employment and rising nominal house prices.
4. A replication of a 'Hong Kong' experience. This was a case where a sharp fall in house prices did not lead to substantial losses in housing lending portfolios, in part, because it was not accompanied by a sustained interest rate increase and in part because of conservative prudential limits on loan to valuation ratios.
5. A replication of three Australian housing loss events over 1984-1994

9.1 'Normal' and 'good time' simulations

The purpose of these simulations is to see whether TUI will replicate the low housing losses that New Zealand banks have normally experienced.

Two scenarios were run. The first is a 'normal' scenario where there is no change in interest rates and unemployment and house prices increase by the mean expected amount (7.5 percent over the three years).

The second scenario has the same no change assumptions for interest rates and unemployment but assumes there has been a housing boom and house prices have increased by 30 percent over the three years.

Estimated PDs for the range of LVR tranches are presented below together with some actual data showing write-off rates for a New Zealand bank over the period 1997-2005. The housing boom scenario shows a very good fit with the New Zealand bank data.

The key point to draw from this analysis is that very low 'benign times' default rates can be generated by a model which has a very much larger long run average default rate. For the 80 percent LVR tranche, for example, the average PD is around twenty times higher than the boom times PD.

Table 33
PDs in Normal and Good times
Annual averages percent

	LVR	*100*	*90*	*80*	*70*	*60*
'Normal' scenario		*0.014*	*0.012*	*0.006*	*0.003*	*0.0003*
'Housing boom' scenario		*0.006*	*0.003*	*0.0015*	*0.0006*	-
NZ bank		-	*.004*	*.0025*	*.0005*	-

9.2 New Zealand FSAP stress test replication

In 2003 the Reserve Bank produced a set of stress tests as part of the FSAP exercise. This involved the Reserve Bank setting out a set of stress scenarios and the four major banks using their own models, and/or, their best judgements to calculate, amongst other variables, the impact of the stress events on their bank's profits for each scenario.

Two of the tests provided information on the banks' housing portfolios. The first was a single factor test. Banks were required to calculate the losses in housing portfolios due to a stress event which combined a 20 percent drop in house prices and an increase in the unemployment rate from 5 to 9 percent. The average credit losses were about 1.1 percent of banks' residential mortgage assets.

The second scenario was a more complex event which modelled a shock to New Zealand's external credit rating. This 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. These variables had a defined path over the three year duration of the shock. The aggregate housing credit losses for the scenario was about 1.0 percent of total housing loans.

The DSR and LVR structure of the New Zealand banking system residential loan portfolio assumed for this exercise was generated by the Bank's Data lab exercise. This provided a portfolio structure as at 2004.

The 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 reflects banks' assessment 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. After the FSAP stress test process was completed it was felt that the house price fall assumption was perhaps too light and it is useful to get a sense of how much difference a bigger decrease in house prices would have made. We found that a decline of 25 per cent increased the loss rate from 1.2 per cent to 2.1 percent.

Finally, the loss distribution generated by the model was used to put a likelihood on the sizes of the losses generated by the stress scenarios. Table 36 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'.

This last exercise is also useful in showing how scenario analysis could be used to test, and perhaps improve, the calibration of the model. The likelihood results show that the loss associated with the 25 per cent house price fall is less than half as probable as the loss with a

15 per cent shock. This result does not seem very intuitive. If house prices can drop by 15 per cent, it might not be that much less likely that they could drop by 25 per cent – given the possibility that a 'snowball' effect could develop once a large interest rate driven decline got underway. Put another way, our result is mainly a function of two model inputs – a normal house price distribution and the relatively low absolute value of the house price/interest rate correlation – which, together, do not generate a sufficiently 'thick' tail. 'Thickening' the tail by changing the interest rate/house price correlation to -0.5 and increasing house price volatility to 20 per cent, changes the likelihood of the 15 and 25 per cent house price shocks to 3.7 and 6.5 per cent respectively, which seems a more plausible estimate of their relative likelihoods.

Table 34
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*

9.3 New Zealand 1980s experience

An example of where housing loan portfolios were robust to a seemingly severe interest rate shock is the New Zealand experience of the mid 1980s. A tight monetary policy saw average interest rates increase by an average of around 50 percent above 1984 levels over the 1986-88 period. While many homeowners came under severe debt servicing stress, bank lending losses appear to have been, on average, relatively moderate. Table 35 shows the PD outputs for a scenario that tries to replicate the 1980s experience. The unemployment change was set at zero (the actual average change was slightly negative) and the average house price increase at 15 percent. The interest rate increase was assumed to be 6 percentage points. New Zealand banks' mortgage portfolios would have been almost entirely concentrated in the low risk buckets in the south east quadrant of table due to the their historically conservative lending policies and to the effect of inflation in reducing LVRs and DSRs after the origination date. It is easy to see, therefore, why losses would have been low in the face of a sharp mortgage lending rate increase. The great bulk of loans would have had an LVR of under 60 percent when the interest rate shock hit and DSRs would mostly have been under 25 percent. The annual loss rates in the most populous buckets would therefore have been under 0.01 percent.

Table 35
New Zealand mid 1980s Scenario
3 yr. PDs by risk buckets

LVR / ***DSR***	***100***	***90***	***80***	***70***	***60***	***50***
40	*3.92*	*1.84*	*1.09*	*0.39*	*0.19*	*0.05*
30	*2.25*	*1.13*	*0.54*	*0.28*	*0.07*	*0.04*
25	*1.85*	*0.91*	*0.48*	*0.19*	*0.06*	*0.03*
20	*1.28*	*0.71*	*0.33*	*0.12*	*0.06*	*0.03*
15	*0.93*	*0.45*	*0.19*	*0.08*	*0.05*	*0.02*

9.4 A 'Hong Kong' scenario

After the Asian crisis in 1998 Hong Kong experienced an extremely sharp fall in residential housing prices. From the peak to the bottom of the house price cycle the fall was in the order of 70 percent. Losses on housing loans, however, do not appear to have been high. We do not have access to actual default or loss rates but the increase in the delinquency ratio (the stock of loans past 90 days due) prior to the crisis (0.3 percent in June 1998) to a peak of around 1.5 percent does not suggest extreme losses. The Hong Hong experience is sometimes cited as evidence of both the resilience of residential mortgage portfolios even to extreme shocks in house prices and additionally, or alternatively, to the success of regulatory policies in ensuring that banks did not engage in excessively risky mortgage lending. In Hong Kong at the time banks were not allowed to lend more than 70 percent of the value of a residential dwelling.

Rather less attention has been made to the role played by interest rates in moderating losses. While mortgage interest rates did respond to the sharp increase in wholesale interest rates during the Asian crisis the increase was muted and short lived. Over the first three to four years of the event there was no material change in interest rates on average.

While we do not have enough information, in particular, about the structure of Hong Kong banks' mortgage portfolios, to model the Hong Kong crisis with any precision there is still some value in modelling a stylised version of the events that unfolded there. An important part of the analysis are our assumptions about the structure of an illustrative bank's mortgage portfolio. Because there had been a rapid increase in dwelling prices in the years preceding the crash the time profile of mortgage originations will have a marked effect on the structure of LVRs just prior to the price crash. We have assumed that the bank had built a portfolio of loans evenly over the seven years preceding the crash and that the average size of its loans was a function of the average current dwelling price. Consequently its portfolio, in value terms, is heavily weighted to more recent originations.

While the peak to trough change in Hong Kong dwelling prices over the full cycle was nearly 70 percent this fall took six years. The fall over the first three years of the shock was around 50 percent and we have used this change in our scenario. We have captured the adverse macro scenario with a 3 percentage point increase in the unemployment rate but assumed no change in the average interest rate over the period. Table 39 shows the cumulative default rates and losses for this base run over the three years by LVR at origination. The results of a second scenario that assumed the above inputs plus a sustained increase in interest rates of three percent points over the scenario horizon are also presented.

Table 36
'Hong Kong'
Portfolio composition by year of origination

Origination Yr.	% share
1991	*6.4*
1992	*10.0*
1993	*12.0*
1994	*16.4*
1995	*14.4*
1996	*16.9*
1997	*23.7*

Table 37
Default rates and losses by LVR at origination

LVR	***90***	***80***	***70***	***60***	***50***
3Yr. PD % Base run	*3.3*	*2.9*	*2.1*	*1.4*	*0.8*
3Yr. PD % Interest rate increase	*6.6*	*5.8*	*4.2*	*2.7*	*1.7*
Loss rate % Base run	*1.6*	*1.3*	*0.9*	*0.6*	*0.3*
Loss rate % Interest rate increase	*3.2*	*2.6*	*1.8*	*1.1*	*0.6*

Assuming that the portfolio was evenly split between the 70, 60 and 50 LVR buckets at origination then the average three year default rate in the base scenario would have been 1.4 percent. This generates an economic loss rate for the three years of about 0.6 percent (and a loss of principal of about 0.4 percent) which is consistent with the observed capacity of Hong Kong banks to take the downturn in their stride.

Our analysis also suggests that the impact of the 70 percent limit on bank's losses may have been less than is sometimes suggested. The 90 percent LVR bucket loss rate was clearly higher than the 70 percent LVR rate but the difference - 1.6 compared to 0.9 percent - was not so pronounced that there would have been a material difference to the overall outcome if the regime had been more liberal. For example if 20 percent of the portfolio had consisted of 90 percent LVR loans at origination then the portfolio loss rate would have increased to 0.8 percent. If, on the other hand, the dwelling downturn had been accompanied by the three percentage point increase in mortgage interest rates then losses would have doubled.

9.5 Australian Housing Losses of the 1980s and 90s

Finally we have simulated the Australian housing mortgage default experiences of the 1980s and 1990s because there is some data covering this period which is sourced from mortgage insurer's loss experiences. Unfortunately the data shows the default experience by the policy year. That is it captures the defaults for all of the years including and subsequent to the year the policy was written rather than by the year in which the default occurred. This makes it difficult to relate default events to the drivers of default but the data can be useful if we focus on the peak loss years. If it is assumed that most of the loss on the policies written in those years occurred over the subsequent two years then we can relate those losses to changes in the macro risk drivers over that period. Measured by the peaks in the data there were three stress events. Table 38 sets out the relevant macro data and the associated default data which is differentiated by LVR. Note that the housing index data refers the greatest cumulative loss over the period subsequent to the peak loss year.

Table 38
Stylised Australian macro scenario and housing default rate data

Shock (peak yr)	***1985***	***1989***	***1994***
Mortgage interest rate increase two yr. average	*+2.2 percentage points*	*+2.3 percentage points*	*+0.4 percentage points*
Unemployment	*-1 percent*	*-1 percent*	*-1 percent*
House prices	*Aus. 86-87 - 4.2 %*	*Aus. 89-91 - 6.6%*	*Aus -2.1% Melbourne-5%*
LVR	***2Yr. default rate Mortgage Insurance data***		
90+	*4.4*	*4.3*	*3.1*
85-90	*2.2*	*2.8*	*1.9*
80-85	*0.9*	*2.2*	*1.4*

Table 39
TUI Default rate estimates for Australia

LVR	***1985***	***1989***	***1994***
90+	*3.8*	*4.6*	*1.8*
85-90	*2.1*	*2.9*	*0.9*
80-85	*1.5*	*2.0*	*0.6*

10 An assessment

The obvious strengths of TUI both in its capital estimation and simulation modes are its transparency, precision and flexibility. It is possible to apply the model to a range of issues; it is generally pretty clear what is driving the results and we have found it to be a very useful diagnostic tool. From a New Zealand supervisory perspective its capacity to capture those aspects of the New Zealand macro-economic environment which affect key risk drivers is particularly important.

That being said there remains the question of whether the TUI results are sufficiently robust to provide some guidance in the setting of capital ratios. Our sensitivity analysis showed that capital results were particularly sensitive to the macro-economic inputs and that there was a reasonably wide confidence interval around our central risk weight estimates

Our assessment is that these uncertainties do not undermine the value of the model. Uncertainty about the macro-economic environment are a fact of life in risk modelling and they affect any model. The results of more empirically based modelling will often depend critically on whether the chosen estimation period is sufficiently representative of the relevant population of outcomes and judgements have to be made about the application of past experiences to the future. The strength of TUI is that these sensitivities are more transparent and the user is forced to think carefully about them.

Figure 10
Anchoring TUI

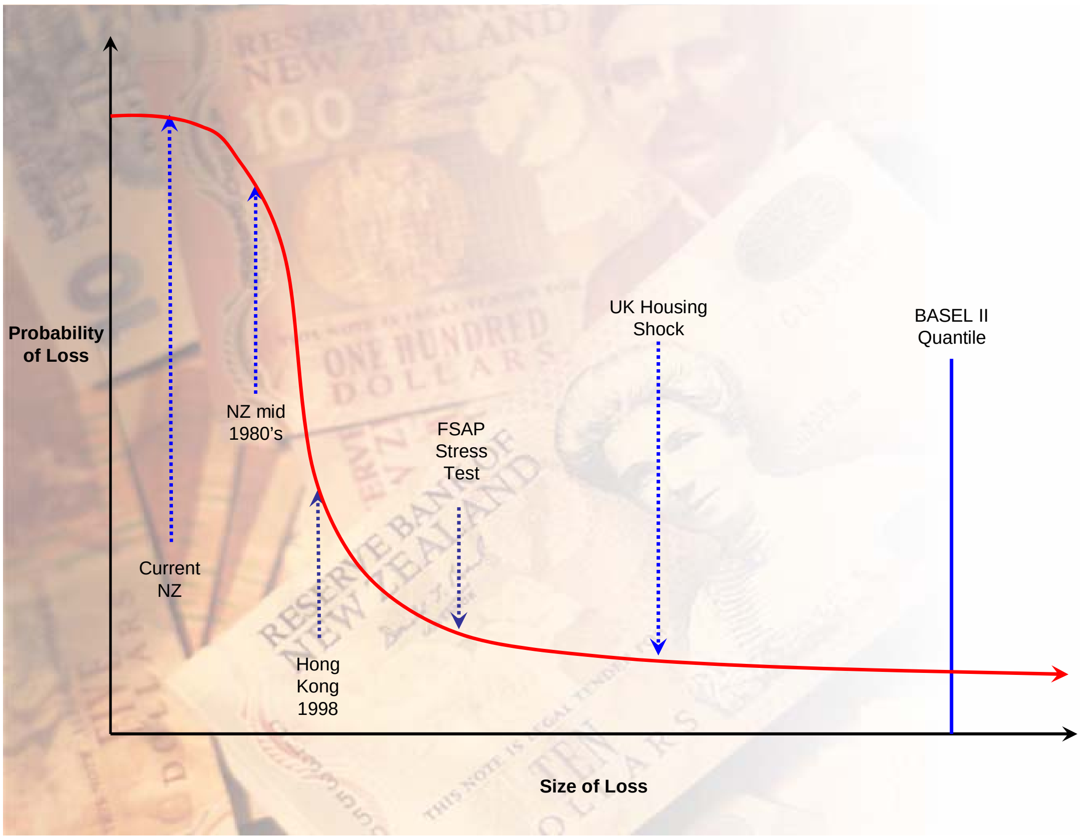

On the behavioural side, on the other hand we believe that TUI is reasonably anchored, at least by the standards that could reasonably apply to tail end event modelling. It has been able to replicate, with a reasonable degree of accuracy: the recent New Zealand ‘good times’ experience; the New Zealand FSAP stress test outputs; the Australian loss experiences of the 1980s and 1990s; and in a more qualitative way the outcomes of the ‘Hong Kong’ house price crash of the 1990s and the New Zealand interest rate shock of the 1980s.

Finally, the current version of TUI is not the last word on this kind of modelling. TUI was conceived as a ‘proof of concept’ model and lends itself to extensions and improvements. On the data side better estimates generated by further research and more recent empirical outcomes can readily be accommodated.

References

Ambrose, B.W. and LaCour-Little, M. (2001) "Prepayment Risk in Adjustable Rate Mortgages Subject to Initial year Discounts: Some New evidence "
Real Estate Economics 29, 305-327.

Basel Committee on Banking Supervision (2004) "International convergence of capital measurement standards; a revised framework".

----------- (2005) "An Explanatory Note on the Basel II IRB Risk weight Functions".

Boheim, R., and Taylor, M.P. (2000) "My Home was my Castle: Eviction and Repossessions in Britain," Journal of Housing Economics 9, 287-319.

Breedon, F.J. and Joyce, M.A.S. (1992) House Prices, Arrears and Possessions," Bank of England Quarterly Bulletin May 1992.

Calem, P. and Follain, J. (2003) "The asset-correlation parameter in Basel II for mortgages on single-family residences," Board of Governors of the Federal Reserve System.

------------ (2004) "Risk based capital requirements for Mortgage loans," Journal of Banking and Finance 28, 647-672.

Calem, P. (2005) "Implementation of Basel II by Mortgage Lenders: Potential Loan Performance Solutions" LoanPerformance

Coleman, A., Esho, N., Sellthurai, I. and Thavabatan, N. (2005) "Stress Testing Housing loan portfolios: a regulatory case Study," APRA Paper presented at the Basel Committee on Banking Supervision conference Vienna.

De Franco, R. (2002) "Unifying Models of severity on defaulted mortgages," Economics Department University of California at Berkley.

Fitch Ratings (2007) "UK Residential Mortgage Default Model Criteria".

Helbling, T. and Terrones, M. (2003) "Real and financial effects of bursting asset price bubbles," in World Economic outlook, chapter II, When Bubbles Burst, IMF, 61-94.

Helbling, T. (2004) "Housing price bubbles- a tale based on housing price booms and busts," BIS papers No. 21.

Kaskowitz, D, Kipkalov, A., Lundstedt, K. and Mingo, J. (2002) "Best Practices in Mortgage Default Risk Measurement and Economic Capital," Loan Performance, Washington Mutual Inc.

Jaffee, D.W. and Welke, G.M (2003) "The Risk-Based Capital test for Fannie Mae and Freddie Mac," Office of Federal Housing Enterprise Oversight.

Lambrech, B., Perraudin, W. and Satchel, S. 1997 "Time to default in the UK mortgage market," Economic Modelling vol. 14, 485-99.

May, O. and Tudela, M. (2005) “When is mortgage indebtedness a financial burden to British households? A dynamic probit approach,” Bank of England Working Paper no.277.

Pain, D. (2003) “The provisioning experience of major U.K. banks; a small panel investigation” Bank of England working paper no.177.

Pennington, Cross A. (2003) “Subprime and Prime Mortgages: Loss Distributions,” OFHEO Working papers 03-1.

Quercia, R.G. and. Stegman, M.A “Residential Mortgage Default: A Review of the Literature,” Journal of Housing Research Vol. 3 issue 2, 341-379.

Rindaldi, L. and Sanchis-Arellano, A. (2006) “Household debt sustainability; what explains household non-performing loans,” European Central Bank Working paper series no. 570.

Whitely, J. Windram, R. and Cox P. (2004) “An empirical model of household arrears” Bank of England Working paper no. 214 2004.

Wong, J., Fung L., Fong T. and Sze A. (2004) “Residential mortgage default risk and the loan-to–value ratio”, Hong Kong Monetary Authority Quarterly Bulletin December, 33-45.

Appendix A

The relationship between average LGDs and LVRs

The historical evidence does not show a positive relationship between LVRs and LGDs. The mortgage industry data shows a long-run average claim of approximately 20 percent and APRA analysis showed that, after controlling for loan size that LGD was negatively related to LVR. This is consistent with Tui outputs which typically show higher LGDs for the lowest LVR risk buckets.

The empiric and TUI results do not accord with the commonly held intuition that the relationship should be positive. This intuition is probably driven by a sense of the numbers which might driven by the following simple exercise in mental arithmetic. If there is a house price shock of, say, 30 percent then the loss on loan A which had an LVR of 90 percent will be 20 but it will be only 5 on loan B which had an LVR of 75. Loan A's LGD should, therefore, be four times greater than that of loan B's.

There are three reasons was this intuitive reasoning is wrong.

1. The idiosyncratic house price effect
2. The cost of default effect
3. The denominator effect

1. Idiosyncratic house price effect
   The impact of idiosyncratic changes in house prices varies by LVR bucket. It has little if any impact on the higher LVR buckets but accounts for all of the failures for the lowest bucket. An illustration of the impact of three representative buckets is shown in table one.

Table 1 LGD and LVR relationships

LVR	*Shortfall on index decline*	*Shortfall after idiosyncratic effect*	*Shortfall after fixed cost of 30 percent*	*LGD as percentage of loan*
90	20	20	50	55.6
75	5	8	38	50.7
60	-	5	35	58.3

2. The fixed cost of default effect
   If a borrower defaults a number of costs are incurred which are invariant to the size of the shortfall between the value of the loan and the value of the house that leads to the default. These 'fixed costs' relate to:

   a. The time value of money
   b. The costs of foreclosure and disposal of the asset
   c. The discount on the value of the house caused by the foreclosure.

   In TUI these factors can total 30 percent or more.

3. The denominator effect

The overall loss should be divided by the value of the loan. Other things being equal this generates a higher proportionate loss for the low LVR loans.

The impact of all three of these factors generates the negative relationship between LVR and the LGD for the 75 and 60 percent risk buckets shown in table one.