

An Appraisal of the RBNZ Proposals on Bank Capital

David Miles

Imperial College

London

August 30th 2019

The RBNZ proposals are that banks in New Zealand should operate with a very substantially higher level of equity funding than is required today. Under those proposed new rules the largest banks would be required to have a Tier 1 capital ratio of 16%, made up of a 6% Tier 1 requirement and a 10% buffer that comprises a 7.5% capital conservation buffer, a 1% large-bank buffer and a 1.5% counter cyclical buffer. Under current rules the Tier 1 requirement is 6% (which would remain) and the overall size of the buffer that sits on top of that is considerably smaller. The buffer on top of the 6% requirement is currently 2.5% and this would rise to 10% for large banks (9% for smaller banks) under the RBNZ proposals *and* it would all have to be in the form of CET1 capital.

Not surprisingly this is controversial. It has generated a good deal of criticism of the approach taken by the central bank. This paper considers the analysis undertaken by the RBNZ and assesses whether it justifies the proposals on bank capital.

1. The approach taken by RBNZ

The approach taken by the RBNZ appears to be to set an acceptable limit on the risk of widespread bank failure and then to assess what level of capital is consistent with that. The wider implications of that level of capital – for example on aggregate incomes – appear to be secondary. This approach seems to be lexicographic: decide on an acceptable amount of banking sector risk to take and then assess whether going beyond that (in terms of capital required of banks) is justifiable because it might generate higher average output.

In fact the idea underlying the analysis that underpins the proposals is much more nuanced than this and does not start from a (somewhat

arbitrary) decision that once in 200 years is the right average frequency to accept the arrival of banking crises. Instead there is more of an iterative process: start with a lower acceptable level of safety (so that crises happen once in 100 years) and then assess whether one could do better in terms of average economic output with higher capital. That is found to be the case (conditional on the calibration of banking risks and of their economic implications). Only at an acceptable risk of banking crises that makes it a 1 in 200 year event is there no clear scope for raising average GDP by going for higher bank capital. So in fact the 1 in 200 year outcome is not a pre-set target; it is more a consequence of the iterative procedure which takes into account the wider economic implications of setting bank capital.

That calculation is of course sensitive to a wide range of assumptions – and the choices made here by the RBNZ in its analysis have generated a good deal of criticism; specifically that they have consistently erred on the pessimistic (or conservative) side of central estimates. I consider the evidence on this in detail below. But one over-arching criticism of the approach taken by the RBNZ does melt away in light of the procedure actually followed by the central bank. This is the criticism that in starting with an (arbitrary) assumption on risk appetite - and not undertaking any cost benefit analysis - the validity of the proposals is undermined. In fact the way that the analysis converged on the 1 in 200 year estimate of an acceptable degree of banking risk was by reference to what that implied about the average level of GDP it generated. And in calculating the impact upon average GDP the analysis was undertaking a cost benefit analysis in which the measure of welfare was average total incomes in New Zealand. That measure of welfare assumes risk neutrality. So rather than being based on an extreme and arbitrary view of risk which could be seen as being excessively averse to low probability bad outcomes, the analysis actually rests on something closer to risk neutrality.

2. The calibration of risk and of what counts as capital.

Key assumptions that drive the RBNZ analysis concern:

- i. The risks of losses facing banks
- ii. The effect of more equity funding – and the type of equity allowed - on the probability of bank insolvency
- iii. The impact of more equity on the weighted cost of bank funds

- iv. The effect of a higher cost of bank funding on lending rates and of the latter upon economic activity

I consider each in turn

- i. The risks of losses facing banks*

In the RBNZ analysis it is the scale of unexpected bank losses relative to the capital of the banking system that is taken to be the fundamental determinant of whether there is a banking crisis. A value at risk model – the asymptotic single factor risk (ASRF) model – is used to explore the relation between capital and unexpected losses. This model has been widely used in the assessment of bank risks and in the various iterations of the Basle capital rules. It is the model specified in Basle III for use by internal ratings based (IRB) banks when estimating risk weights. There are three inputs which affect the probability distribution of bank losses in this model:

The unconditional probability of default - **PD**; this is the typical proportion of borrowers that default in a year averaged across all states of the economy (both good and bad).

The correlation of losses across banks – **R**; this measures the extent to which defaults are influenced by a common economic shock rather than by idiosyncratic shocks to specific loans;

Loss given default -**LGD**; this is percentage of a loan exposure that a bank will lose if the borrower defaults.

The product of PD and LGD is an estimate of expected losses which are assumed to be covered by provisions.

The higher are R, PD and LGD the greater capital will be required to maintain a given risk of a banking crisis – that is a situation where bank capital is overwhelmed by the scale of bank losses. The choice of these three parameters has a material impact on the assessed quantity of bank capital required to generate a given level of banking sector risk. Table 2 of the RBNZ Capital Review Background Paper “*An outline of the analysis supporting the risk appetite framework*” (April 2019) shows the sensitivity of required levels of bank capital to the three parameters. A choice of PD of 3% rather than 2.25% raises the estimated required amount of bank capital by around 20%; using a value of R of 0.35 and of LGD of 45% - as compared with base case values of R=0.3 and LGD = 0.4 - raises the required amount of capital by almost 30%.

The RBNZ considered a plausible range of values for the three parameters. The ranges, and the central values used for many of the simulation exercises (in brackets), are as follows:

PD: 1.5-3.0% (**2.25%**)

R=0.2-0.4 (**0.3**)

LGD=0.35-0.5 (**0.4**)

The RBNZ analysis did not just rely on assessed central values to inform proposals on capital requirements. Monte Carlo simulations were done to assess the range of values for appropriate capital (relative to unweighted assets – i.e. leverage ratios) consistent with the assumed risk tolerance. The results are shown in Figure 9 of the Capital Review background paper. The central estimate (the mode) is around 9% - consistent with a required capital to risk weighted assets ratio of slightly below that in the RBNZ proposals. But the skew of that distribution is to the upside meaning that there is more chance that the appropriate level of capital is well above the modal outcome than well below it.

Using the RBNZ central values for the three parameters generates a required level of capital – given the assumed risk tolerance – in line with its proposals.

In submissions many banks, and banking associations, have argued that the central values used for the key parameters are unduly pessimistic on the scale of risks of losses on loans. Recent experience of non-performing loans and of losses would suggest that a figure of 2.25% for average probability of default, and of 0.4 for loss given default, are relatively high. New Zealand has not experienced a banking crisis in recent history and avoided serious damage to its banking sector in the 2007-2008 global financial crash. Whether or not that reveals an enduring level of above average banking sector stability or a period of above average luck is hard to judge. If one took the experience of all countries with levels of per capita income comparable to New Zealand and assessed their recent banking sector losses the central values used by the RBNZ do not look at all pessimistic; if one considered the experience of those other countries as uninformative and used only recent New Zealand history as a guide the figures do look pessimistic about risks of losses.

Evidence on appropriate values of R (correlations across bank losses) for New Zealand is thin. It is likely that such correlations are higher in a period of general economic stress¹ ; this means that the absence of recent sharp shocks to

the New Zealand economy makes accurate assessment of appropriate values for R a challenge. Experience in other countries would suggest a value of around 0.2 – at the low end of the range used by the RBNZ and materially below the central value of 0.3 One justification for using a higher value than for other countries is that house prices (and some other asset values) seem more highly correlated with GDP in New Zealand. But this evidence is quite weak and not a firm basis on which to be confident that a higher value of R is justified.

Overall judgement – the RBNZ base case (or central values) of the key parameters could be argued to err on the side of caution in assuming risks of bank losses in New Zealand are no lower than in other countries whose recent experience has been of materially higher losses. Cumulated across the several relevant dimensions (R, LGD, PD) this might bake into the analysis a degree of caution relative to a neutral assessment of the factors which treats risks as symmetric. Yet in no case is the divergence from a central estimate clear cut – and the RBNZ analysis recognises where the assumed value for NZ are somewhat different from those based on data from other countries and sets out reasons for such divergence.

This degree of caution in some parameters (perhaps most clear cut in the choice of R and, to a lesser extent, in PD) should be balanced against other areas where the RBNZ analysis has chosen values on the less cautious side of a central value. The evidence that there has been a degree of excess optimism - in particular as regards the costs of a banking crisis - to offset the tendency towards cautious pessimism in the parameters described above is considered in more detail below.

ii. The effect of more equity funding – and the type of equity allowed - on the probability of bank insolvency

Two key judgements made by the RBNZ in assessing the impact of bank capital on the risks of effective bank insolvency are these: first, that the point at which a crisis occurs is when the level of capital is reduced to zero; second, that capital is overwhelmingly in the form of CET1 – effectively this means equity funding.

The decision in the analytical work that underpins the RBNZ proposals to take complete exhaustion of capital as the point of crisis errs on the side of permissiveness; it suggests a less conservative, or less averse, attitude toward

¹ See, for example, “Basle III Correlations”, Special report by Fitch Ratings, November 2011.

risk than might have been taken. This is because in practice a bank with CET1 relative to assets of only a percent or two is likely to be in a crisis and find it hard to get funding. In contrast the decision to exclude pretty much everything besides equity from the required amount of capital – which is certainly not without some justification – could be seen to err the other way; it has certainly generated criticism from banks who consider it substantially increases the cost of complying with higher capital standards.

Under the RBNZ proposals only common equity and preference shares that are non-redeemable (with no contingent triggers) should qualify as Tier 1 capital; *all* of the 10% (9% for smaller banks) prudential capital buffer will be in the form of CET1 capital (equity). CET1 capital unambiguously absorbs losses on a going concern basis over a long time horizon. Contingent debt – debt that is written off or converts to equity when triggers are pulled – will not count.

In principle contingent debt would seem an attractive form of bank funding to help preserve financial stability because triggers can be set at levels well clear of bank failure so that contingent debt could boost shareholder capital when the bank was still reasonably well capitalised. In practice there is limited evidence of smooth conversion at points well before insolvency which would help to maintain a bank as a going concern. Convertible instruments that would meet current requirements to count as AT1 or AT2 are legally quite complex. Nor is the theoretical case that contingent debt improves bank incentives on risk-taking clear cut; there are cases where in theory contingent debt encourages excess risk taking by bank managers acting in the interests of shareholders. Setting appropriate triggers is also challenging: using accounting measures of equity means that the measure to compare with the trigger threshold may be out of date – a potentially serious issue in a fast moving situation. But using stock market values means conversion can be triggered as a result either of a panic or of a downward spiral in market values itself created by the prospect of a critical threshold being crossed.

Despite all this contingent debt seems very attractive to banks who argue that it is an effective way of meeting capital requirements. This is largely on cost grounds. Some banks put the cost of Tier 2 capital and senior bail-in debt at between 3% and 3.6% while the cost of equity capital is put at around 14.5%. The reasons for the apparent enormous difference in cost are something of a mystery. Standard textbook corporate finance would say that there should be no overall cost advantage over equity (the Modigliani Miller result). Tax advantages of debt over equity are one reason why in practice there could be an

advantage of contingent debt. But the imputation system in New Zealand (and in Australia) means that this source of advantage should be low.

Perhaps the apparent huge cost advantage might reflect unrealistically low perceived chances of conversion triggers being pulled. Perhaps a belief in very low chances of conversion or write-off of contingent debt reflects a rational belief that the authorities would be wary to trigger conversion or write-off of hybrid capital instruments. Neither interpretation is encouraging for the view that such contingent debt can be relied on to maintain stability.

A key point is that the RBNZ is focused on reducing substantially the probability of insolvency. Contingent debt is a form of loss absorbing capital – it can protect large classes of other debt holders (and the NZ government) from losses should bank insolvency become very likely or inevitable. Whether it is effective in reliably reducing the probability of bank insolvency and allowing banks to remain going concerns is less clear.

Two things would strengthen the case for allowing banks to use contingent debt as part of the prudential capital buffer:

- More evidence on the effectiveness of contingent debt in preserving banks as going concerns in the face of significant losses
- A clearer picture of why contingent debt seems to be much cheaper than equity

It would not be wise for a government to increase speed limits on roads in the anticipation that a relatively new, untested and remarkably cheap – but potentially more effective - type of air bag was to be used by car manufacturers. Like all analogies this one is imperfect. But it suggests to me that the proposal to exclude contingent debt, at least initially, from the capital buffers required of banks has merit.

iii. The impact of more equity on the weighted cost of bank funds

The conceptual approach taken by the RBNZ to translate a given rise in the use of equity funding through to the impact on the weighted average cost of bank funds is standard; it is neither complicated nor particularly controversial. The impact of using more equity depends upon the extent to which the providers of equity require a higher average return than those who provide debt and also on whether that risk premium might decline if there is more equity, which reduces the variability of returns. It also depends on whether equity returns are treated less favourably for tax purposes – for example if debt servicing costs are tax

deductible while returns to shareholders are paid out of post-tax income. (The precise formula used to calculate the overall effects is equation 3, page 36 of RBNZ paper “*An outline of the analysis supporting the risk appetite framework*”, Susan Guthrie, April 2019).

While the theory behind the approach used to estimate the impact on the average cost of bank funds of using more equity is clear enough, the empirical application of the approach is more tricky. Three issues are of first order importance: the required cost of equity funding relative to debt; its sensitivity to changes in the balance of equity to debt; and the tax implications of a change in bank leverage.

RBNZ assumes a cost of equity funding to banks in NZ of slightly above 14%. That seems high – particularly to the extent that it relies on past returns on the historic cost of bank equity (that is the book equity) rather than its market value. The largest banks that operate in New Zealand trade at a price to book premium so that the return on the market value of equity would be significantly lower than the 14-15% figure and much closer to 10%. The 14% figure compares to an assumed annual average cost of debt of 2.3%, generating an equity risk premium of just over 11.5% - a value substantially higher than is considered likely for an overall equity risk premium on all stocks which is closer to 5% (and which is rarely considered to be as high as even 7 or 8%).

How much lower the expected return might be if equity risk was lowered because gearing was reduced is even harder to judge. RBNZ use a Modigliani-Miller (MM) offset of 50%. This is a material factor because should that offset due to lowered equity risk be worth only 25% rather than 50% the estimate of how much greater the cost of funds is as equity rises would be almost 50% higher; if the offset was 75% the impact of the cost of funds would be about 40% lower. The estimated GDP cost of the extra cost of bank funds would also be either 50% higher or 40% lower relative to the central case when the MM offset is assumed to be 50%.

The 50% MM offset is not out of line with findings from other countries. But there is one reason why experience from other countries may be somewhat misleading. This is because of the unusual position of NZ where the 4 dominant banks (whose total market share exceeds 80%) are all wholly owned subsidiaries of Australian parent banks. It is the Australian parents that would be issuing new equity if that was the preferred route to satisfy higher capital ratios in NZ. The case has been made strongly by the Australian banks that shareholders buying new equity that they (and not the NZ subsidiary) issue are unlikely to demand a lower rate of return because the (relatively small)

subsidiary has lower gearing. Whether they would or not is hard to judge; if there was no offset for lower risk it would imply a degree of irrationality or myopia on the part of investors in equity or perhaps a failure of Australia parent to successfully explain that the extra equity would reduce risk in the part of the banking group where the funds would be directed. And even if there was no reduction in the required return of shareholders that does not prevent the Australian parent taking into account the lower risk in its internal setting of hurdle rates on capital made available to its NZ subsidiary.

If the RBNZ assumption on the MM offset might have erred on the side of reducing somewhat the impact of using more equity on the weighted average cost of bank funds, the tax assumption looks likely to have gone in the other direction. The assumption here seems to assume the classic tax distortion that favours debt over equity: debt servicing costs are tax deductible while the cost of servicing equity through paying dividends are not. That is why in the formula used by the RBNZ the weighted cost of bank funds rises by the product of the corporate tax rate and the interest cost of debt when less tax efficient equity is substituted for debt. But this would seem to allow no role for the operation of the imputation system – which operates in both New Zealand and in Australia – which is explicitly designed to remove (or at least reduce) the tax distortion in favour of debt.

These three factors pull in different directions – the operation of the imputation system and the assumed high required equity premium might be pushing the RBNZ estimates of the cost impact on banks of using more equity above a central estimate; the MM offset may be pushing it the other way. It is possible these factors offset each other leaving the RBNZ central estimate at a plausible level. If, for example, the MM offset was at 25% (rather than 50%), the bank cost of equity was 11% and not 14% (generating a bank equity risk premium of about 8.5%) and the tax effect was absent (because of the imputation system) the estimate would be very close to that assumed by the RBNZ. But this is not a strong justification of the RBNZ figure and one would want to explore the sensitivity of the overall judgement on the appropriate amount of capital given the uncertainty of the right level for the MM and tax offsets and the true cost of bank equity. I consider this in more detail below.

iv. *The effect of a higher cost of bank funding on lending rates and of the latter upon economic activity*

In the RBNZ analysis of the impact of banks using more equity it is assumed that the rise in the cost of bank funding with higher capital requirements is matched by a combination of a rise in interest rates on bank lending and a decline in the return on bank debt. But for the purposes of the estimation of the impact on GDP it is taken that all of the rise in funding costs feeds through into higher lending rates. This seems an appropriate assumption.

How much such a rise in bank lending rates affects GDP is not easy to gauge based solely on NZ evidence. So, as in many other cases, the analysis by the RBNZ relies on average estimates made for a group of developed economies (drawing on research from the US Federal Reserve, the Bank of England, the BIS and an earlier 2016 RBNZ study). A value of unity is used by the RBNZ as the central estimate of the percentage change in underlying GDP for a given change in the lending rate brought about by a requirement for higher equity (lower leverage). So if the lending rate were to rise by 10bp the level of GDP would be lower by 10bp. In a study of 2013 I estimated that an 18bp rise in bank funding costs would generate a fall in UK GDP of about 15bp implying a slightly below unity ratio of the percentage change in GDP to a change in bank lending rates². US evidence is more in line with the RBNZ assumption³.

In various submissions New Zealand banks have argued that the GDP cost of a given rise in bank lending rates is higher than is assumed in the central case by the RBNZ. The higher end of the range of sensitivities given by the RBNZ is 1.2 – still some way below the scale of impacts implied by some of the bank submissions to the Capital Review.

It is important to note that the sensitivity of GDP to banks use of more equity is a function of the product of the effect of higher equity use on banks funding costs and the impact of higher lending rates on GDP. If the RBNZ have overestimated the first effect because of using too high an assumed required rate of return on equity then that would offset a possible under-estimate of the impact of a change in lending rates on GDP (albeit one that is in line with international studies). A quick sense test is useful: if instead of using 14% as the required return on equity the RBNZ had used 10% this would lead to the same overall conclusion on the GDP cost of higher bank equity as if the sensitivity of

² "Optimal Bank Capital", Miles, Yang and Marcheggiano, The Economic Journal, 2013.

³ See Firestone, and Ben (2017) "An empirical economic assessment of the costs and benefits of bank capital in the US", Finance and Economics Discussion Series, 2017-034. Board of Governors of the Federal Reserve System.

GDP to bank lending rates were around 1.45 rather than the central estimate of unity. That would be above most international estimates of the impact of higher bank lending rates on economic activity.

While relying on international studies for assessing the sensitivity of GDP to a higher cost of bank lending is sensible it obviously raises an issue of whether the sensitivity in New Zealand might be different. As in other places in the RBNZ analysis, it is not clear whether this reliance on data from outside NZ biases the results one way or the other. The case for using evidence only from NZ is weakened by the fact that it much limits the historical evidence that can be brought to bear. But it also means there is a question about how relevant average international experience is to the case of NZ; once again that reinforces the case for testing the robustness of the policy recommendations to variation in a range of parameters – an issue I return to below.

3. The limited use of IRB and the need to have buffers on buffers

The RBNZ proposals would significantly reduce the gap between assessment of risk weighted assets based on the standardised approach and the internal ratings based approach used by the 4 major banks. The combination of floors on the ratio of IRB to standardised risk weighted assets (at 85%), and the application of a 1.2 scalar to IRB levels, would mean that IRB levels of assets would probably rise to around 90% (and perhaps higher in the case of some assets) of the standardised levels. Not surprisingly the larger banks consider this proposal to be excessive and particularly dis-advantageous to them.

Two general points are worth making.

First, the aim of reducing reliance on internal models, which are both difficult to assess and which can generate puzzling differences in risk assets for banks which appear to hold similar portfolios, is not unique to the RBNZ. Many central banks and other supervisors share the uneasiness about the IRB approach which seem to generate some surprisingly low assessments of risk weighted assets for some banks. Furthermore, the limited resources of the RBNZ means that understanding how the IRB models have been applied is particularly challenging. In principle it makes sense that banks should use their detailed knowledge of their current and past portfolios – and the credit losses they have experienced – to generate a more accurate assessment of risk than the broader brush treatment under a standardised approach. But if the central bank has limited resources to assess the way banks use their detailed information to

model risk – and if those banks face incentives to reduce their required capital⁴ – there is an obvious tension with the IRB approach.

Second, any greater allowance for the IRB approach which reduced RWA could – given the overall assessment of the appropriate level of bank capital – just be offset by a higher implied need for CET1 relative to RWA. This is plausible because the RBNZ proceeds by assessing how much capital relative to total assets is needed and then converts that into a CET1 to RWA ratio by applying an assumed ratio of RWA to total assets. If RWA to total assets falls because greater allowance is made to use internal models then, absent any changes in assumptions, the implied appropriate level of CET1 to RWA would rise.

Whether the RBNZ proposals would remove significant benefits of more precise calibration of capital to the risk of bank portfolios is not clear. It is possible that the output floor (at 85%) would become the binding constraint on the IRB banks in which case differentiation between them in their assessment of risks for capital purposes would be minimal. What is not clear is whether such differences now largely reflect real differences in risk or more the use of different modelling assumptions applied to broadly similar pools of risks.

Banks have raised another concern about the pressure they would face to use equity capital in excess of the – much higher – levels required by the RBNZ. Currently most banks have a level of capital a few percentage points of RWA above the levels prescribed by the RBNZ. If that buffer-on-the-buffer carried through to the new regime the larger banks would have a CET1 to RWA ratio of around 18%. There is a case here that this concern could be much reduced if the RBNZ might make it very clear that a temporary dip in the level of CET1 below the (substantially higher) conservation buffer under the new regime is of significantly less concern than a similar size shortfall under the current regime where the conservation buffer is less than one third as large. The RBNZ proposals do contain some examples of how such a sliding scale of concern might work as capital dips below the 16% (15% for smaller banks) level:

“For example, when Tier 1 capital is between 13.5 percent and 16 percent, discretionary payments such as dividends to shareholders, or distributions on Tier 1-qualifying preference shares will be limited to 60 percent of net earnings. If Tier 1 capital is below 13.5 percent but above 11 percent, discretionary payments will be limited to 30 percent. At Tier 1 levels below this, no discretionary payments will be permitted. “

RBNZ Capital review Paper 4 “How Much Capital is Enough”, January 2019

⁴ The detailed submissions made by banks to the RBNZ capital review makes it abundantly clear that banks believe there are strong incentives for them to reduce required capital.

But this is an illustrative example and not a detailed proposal. The general idea though is clear enough:

“The high level concept is that as banks go further into the prudential capital buffer, the nature and consequences of the supervisory response will increase in severity.... The details of the nature of the supervisory response, the precise trigger points, the timeframes and obligations on banks to respond, and so on are important details, which warrant a standalone consultation with stakeholders”

“How Much Capital is Enough”, RBNZ, January 2019.

If the RBNZ does not consider that banks need have as large buffers of CET1 over and above their proposed levels as banks feel are required at today’s much lower levels of required capital then it could be helpful to make that explicit. In part the operation of the counter-cyclical buffer could be used to achieve that end – with its size varied as conditions change so that the overall level of capital (including buffers on buffers that banks chose to hold over and above required levels) is not an unchanging amount relative to risk weighted assets. In the RBNZ proposals that counter cyclical capital buffer is set at 1.5% in normal times. The idea is that it could be reduced so as to allow banks to better support lending in the recovery from a shock that generated financial losses. So while the CCYB and the escalating regulatory response to banks dipping below required capital levels are clearly distinct concepts they share a common aim which is to allow banks to continue operating at closer to normal levels after a temporary hit to their capital than would be the case with a more rigid system of requirements.

Whether or not the CCYB is the most useful way of building in such flexibility is not quite so clear. It certainly adds a small amount of complexity to the rules. More significantly its effective operation would seem to require that the RBNZ is able to assess where in “the cycle” the economy sits. This is no mean feat. Telling the difference between an economic shock to banks that is complete and one that is merely the first stage of a crisis that is about to become much worse is difficult. In the UK the hit to confidence following the demise of the mortgage lender Northern Rock at the end of 2007 proved not to be the eye of the storm – in fact it was a mild precursor to something that was much more serious and which came in late 2008. A case could be made on grounds of simplicity (and in the light of the difficulty of working out where in the cycle an economy is) to roll the CCYB into other parts of the prudential capital buffer leaving a single buffer of 10% for larger banks and 9% for smaller banks. But in practice what is far more important than how such a prudential capital buffer is comprised is its overall size.

4. The cost of banking crises:

An essential element of the estimation of the costs and benefits of reducing the risks of bank insolvencies is the calibration of the economic cost of banking crises. The 63% number used as the base case for the Net Present Value (NPV) of the GDP costs of a banking crisis is the average used across a range of studies. A case could be made that this figure is an underestimate – and perhaps a serious underestimate - of the impact of a widespread insolvency right across the NZ economy and a near total breakdown in the provision of bank finance (added to which would be major disruption to the payments system).

International evidence on the scale of the fall in GDP since the banking crisis of 2008, its persistence and the low real discount rate that seems appropriate today might put the NPV lost GDP at well over 100% of annual GDP – and possibly nearer 200%. In the UK 10 years on from the banking crisis the level of GDP is not far off 20% lower than a continuation of the trend that it appeared to be on pre 2008. One should be open minded about how much of this enormous GDP shortfall is due to the banking crisis and about how much of it will be permanent. As regards the persistence of the shock, it is now a universal assumption by economic forecasters in the UK that *none* of the lost output since 2008 will be recovered through a period of higher than average growth that is a catch up to an undershoot from the fundamental path for incomes. If one assumed that only one half of the shortfall since 2008 in the level of UK GDP (ie 10%) was due to the banking crisis, and that all of this is permanent, the NPV of the cost is that of 10% of annual incomes discounted from today to the end of time. At a real discount rate of 3% a year that would generate a cost of around 330% of GDP – and that is at a discount rate that is 5% *above* the level of real yields on longer term UK government debt. (The real yield on 20 year UK inflation proof gilts in mid 2019 was around -2%. The yield on 10 year NZ conventional government bonds at the end of August 2019 was only just over 1%; at 25 year maturity the yield was barely 1.5%. At an inflation rate that averaged 2% the implied real yields on NZ government bonds are mildly negative).

Seen in this light the 63% figure used by the RBNZ seems optimistic.

5. The robustness of results:

The RBNZ analysis of the appropriate levels of bank capital is sensitive to a range of factors about which assumptions must be made. A good deal of evidence is presented by the RBNZ to justify its assumptions. Nonetheless there

is, as the RBNZ recognises, considerable uncertainty about several key parameters. In this section I assess how robust are the RBNZ policy proposals to that inevitable uncertainty about the economic environment.

The key uncertainties include those that reflect:

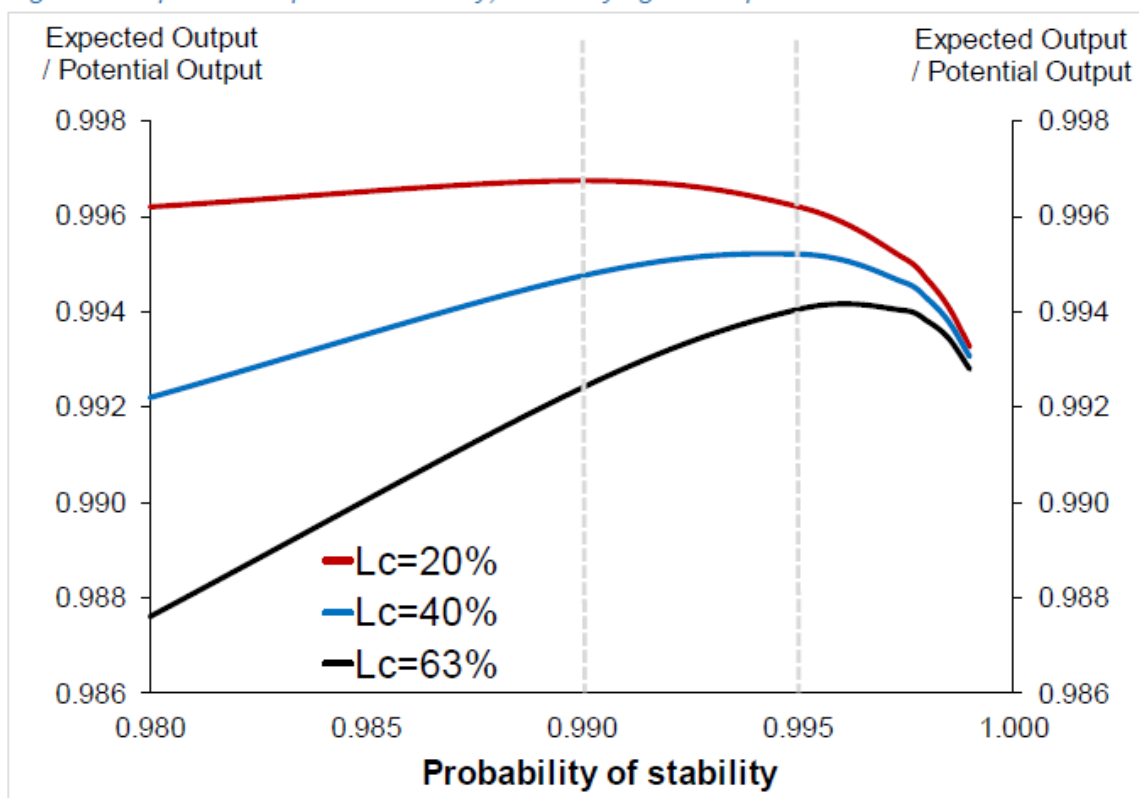
- the riskiness of bank assets (key assumptions here are about R, PD, LGD)
- the effect of changes in leverage on bank average funding costs (key assumptions here are about the cost of equity at current leverage and the size of the MM offset)
- the sensitivity of aggregate incomes (GDP) to changes in the interest rates on bank loans
- the damage done to incomes, now and into the future, by a banking crisis (key assumption here is the NPV of the hit to GDP)

Some sensitivity analysis was included in the background papers made available by the RBNZ. Figure 5 of the April 2019 paper by Susan Guthrie (“*An outline of the analysis supporting the risk appetite framework*”) shows the impact of varying the assumed hit to incomes from a banking crisis. The chart is reproduced below. At the base case assumption (a NPV cost of a banking crises of 63% of current annual GDP) average NZ output over time is maximised at a level of bank capital that implies banking crises have a frequency of slightly less than once every 200 years (ie a probability of a crisis of a bit under 1-0.995). To maximise average incomes one would then need to set capital at slightly above the 16% level (at around 17%) that the RBNZ believes consistent with a chance of bank insolvency of 1-0.995.

The figure shows that if the damage done by a crisis is substantially lower (a NPV of 40%) the optimal frequency of banking crises is close to 0.995.

I argued above that 63% could well be a substantial under-estimate of the cost of a banking crisis. Using a higher figure would of course generate a larger estimate of the appropriate level of bank capital. But other parameters set at base case values may have overestimated risks. To assess the robustness of results it is useful to vary several assumptions (from base case levels) simultaneously.

Figure 5: Expected output and stability, with varying assumptions on cost of crises



Source: “An outline of the analysis supporting the risk appetite framework”, RBNZ, April 2019.

One combination of alternative assumptions that I consider relevant is for a higher cost of banking crises (set at 100% NPV loss in output rather than the 63% base case) but a substantially lower risk of systematic risk of losses on bank assets (setting R at 0.2 rather than the 0.3 base case used by RBNZ and assumed in figure 5). This combination of changes (keeping all other parameters at RBNZ base case levels) implies a level of bank capital (CET1 to RWA) to maximise incomes over the long term of just over 14%.

Just lowering the unconditional expectation of bank defaults from 2.25% to 1.5% - while keeping the NPV of losses from banking crises at 100% of GDP and setting all other parameters at the base case levels – makes the level of CET1 to RWA that maximises average incomes over the long term about 16.5%. This is close to the level in the RBNZ proposals for bank capital.

Reducing the assumed level of losses given default brings the level of bank capital to RWA that maximises average incomes down quite substantially: using an assumed LGD of 0.25 (as against 0.4 in the base case), and sticking with a

100% NPV cost banking crisis, brings the “best” CET1/RWA ratio down to about 13.2%.

Lowering the MM offset has a relatively small impact on the results. Simultaneously setting the MM offset at 0.25 (against 0.5 in the base case) but having the higher NPV of the hit to GDP of a banking crisis (100%) makes the level of CET1/RWA that achieves the highest incomes over time come to be around 18.5%.

It is certainly possible to lower the assumed level of banking sector risks so that it sits within the range of possible plausible values and generates substantially lower levels of bank capital ratios that achieve the highest long run average levels of incomes. If the unconditional probability of default is 1.5% (base case of 2.25%) and the level of exposure to systemic risks is also substantially lower with $R=0.2$ (base case 0.3) then with the NPV cost of banking crises set at 100% the “best” capital ratio falls to just over 11%.

These results show that there are sets of assumptions which are not implausible and which could justify both materially lower *and* higher levels of bank capital relative to RWA than the RBNZ proposals. To my mind they do *not* show that the assumptions the RBNZ made in the base case – viewed in their entirety – have skewed the results towards proposals for bank capital that are clearly higher than they should be.

6. The transition period:

The RBNZ proposals would mean that large and smaller banks operating in New Zealand will need to raise substantial amounts of new equity. Smaller banks – who use the standardised approach to measuring RWA – would need to raise around \$1 billion of equity; the IRB banks (the large 4 Australian-owned banks) would need to raise around \$19 billion. The RBNZ estimate that should banks rely upon retained earnings to build up the extra capital it would take the smaller banks over 7 years to reach the new required level; the 4 larger banks might make the transition in around 5 years. In proposing a transition period of 5 years the RBNZ has set a target that – absent new issues of equity – would probably require larger banks to pay no dividends over the transition and smaller banks might need to shrink their balance sheets. For the larger banks there is the option of the parent raising new equity, which might allow dividends to be paid. Whether that is a more palatable route to compliance is unclear – in some ways it hardly seems much of a gain to raise equity merely to

allow dividends to be paid with no change in the net requirement for extra equity.

For some of the smaller banks that are mutual there is no option of new share issues and so there is a greater chance of them shrinking balance sheets to comply with the proposed rules over a five year transition. But mutuals may be able to issue preference shares. The RBNZ proposals would restrict the use of preference shares to make up no more than 1.5% of risk weighted assets and require that they be irredeemable. That restriction has a logic to it: such shares that were retired in a market where the redemption was taken for granted but where it was problematic to issue new shares would mean a sudden fall in capital. There may be ways of avoiding this if redemption was allowed but subject to restrictions that would prevent a sudden sharp fall in capital. That seems worth exploring, especially if mutuals – with little alternative routes to raise capital quickly - would find it very hard to issue irredeemable instruments.

The key question for the transition path is simple: what is the horizon over which to allow banks to comply with the new rules so as to best trade off the benefits of reaching a more appropriate level of capital quickly against possible costs of a less than very gradual transition? While the question is simple finding a good answer is not – in many ways it is more difficult than the question of what the appropriate end point should be. One answer, which is both simple and has some plausibility, is this: that for those banks for whom issuing new equity is feasible a speed of transition any longer than the period for which retained earnings can generate the extra capital is unnecessary. The logic is this: that any cash flow harm done to owners of equity who were relying on receipt of dividends can be avoided by the bank (or its parent) issuing new shares to those who are not so cash-flow constrained. An alternative interpretation is that for banks with equity traded in a liquid market any cash flow difficulty felt by existing share owners who had expected dividends can be avoided by their selling shares (this is just another Modigliani Miller result).

This would suggest that for the larger banks the 5 year transition path might be appropriate. For the smaller banks – and particularly for the mutuals – things are less easy. The use of near-equity forms of funding for mutuals would help ease the transition.

7. Overall conclusions:

The analysis conducted by the RBNZ seems to have been done with care and in an open minded way. The thoroughness of the analysis is impressive given the relatively small team that worked on it. The details of the analysis have now been made widely available; this has meant that many thoughtful submissions on the RBNZ proposals have been made. Those submissions – and discussions with those who made some of them – have greatly informed my assessment.

The claim that RBNZ analysis consistently errs on the side of caution in modelling risks, and also errs on the side of favouring more equity in its assessment of the impact on the wider economy of banks using more capital, does not seem to me to stand up. While there are some areas where parameter values used are probably on the cautious (or pessimistic) side of a central value (R, PD, LGD) there are other areas where the base case assumptions used in the RBNZ analysis seem to me to be overly optimistic - for example on the scale of the long term negative impacts of a banking crisis. The assumed required return on bank equity also seems to me very high – and this would reduce the estimated level of appropriate bank capital.

The higher level criticism that the RBNZ started with an arbitrary and excessively risk intolerant assumption about acceptable chances of banking sector crises – the 1 in 200 year rule - does not take account of the crucial role played by the estimated impact of bank capital upon average aggregate incomes (i.e. GDP) in New Zealand. The conclusion reached by the RBNZ on acceptable levels of banking sector risk is in fact a result of considering potential trade-offs between greater security and lower average levels of aggregate incomes in New Zealand. Embedded within the analysis that led to the 1 in 200 year rule is a form of cost benefit analysis which uses average GDP as a measure of the relevant wider economic outcomes - a measure which implicitly assumes risk neutrality.

That the RBNZ proposals would make banks in NZ better capitalised than in almost any other developed economy – and is in that sense “out of line” – is in itself not a powerful argument that the proposals go too far. They should be judged in terms of their effects on banks and on the wider economy in NZ – which is what the RBNZ analysis sets out to do. If the answer from a careful analysis is that capital will be different from that in most other countries then so be it; other countries might usefully ask whether the divergences are a reflection of economic differences between NZ and their own country or whether they

have not done enough of the sort of analysis done by RBNZ. The criticism that differences generate a non-level playing field which is harmful is a weak one – do we believe that the right speed limit on roads should be uniform across countries when road conditions (as well as possibly the appetite for risk and the quality of the emergency services) obviously are not uniform? The claim that a “non-level playing field” is a problem is no more persuasive than the claim that “one-size fits all” is optimal.

The focus of the RBNZ analysis is on reducing the chances of bank insolvencies and not on ways of ameliorating the impact of such insolvencies. In taking this route the RBNZ is right to emphasise the damage banking crises can bring – the aftermath of the 2007-08 banking crises in many countries has generated plenty of evidence of this. Thankfully for New Zealand this recent evidence comes from other countries, but their banking systems would not seem to be so different as to make their experience irrelevant. Overall the RBNZ has ended up being quite optimistic in its view about the relatively limited damage done by a banking sector crisis.

Whether or not the RBNZ has come up with a proposal that requires too much or too little capital from its banks is not completely clear cut. That prompts a difficult question about whether the risks of being one side of the right level rather than the other are symmetric. What is clear is that an appropriate level of bank capital cannot be expected to generate zero risk of widespread bank insolvencies. As Larry Summers observed in a different context, if you thought the right way to plan a trip to the airport was to reduce the chances of missing your flight to zero you would spend far too much of your life waiting at airports. But that valid point does not imply that the costs of being a bit too early or a bit too late are symmetric - waiting 10 minutes more than you optimally should at the airport and missing a flight by 10 minutes do not generate the same cost.

One final point leads me to the conclusion that the RBNZ has probably not over-estimated the appropriate level of bank capital. This concerns supervisory philosophy; perhaps strategy is a less pompous term. The RBNZ has adopted a principle of being conservative as regards bank capital to offset possible risks from its light-handed approach to supervision. That is a choice and one partly based on the view that having very large resources devoted to intrusive oversight of banks is not the most efficient road to go down. That is a conclusion that engineers and safety experts often apply when dealing with the

design of structures. There is a choice between building bridges many times stronger than you expect them to need to be OR you having large teams of inspectors who pay frequent visits to examine all bridges and monitor flows of traffic over them. It is clear that nearly all countries follow the first strategy.

That may be a useful guide for bank supervision.