



BANK OF ENGLAND

Speech

The Bank of England's approach to stress testing the UK banking system

Speech given by

Alex Brazier, Executive Director for Financial Stability Strategy and Risk, Member of the Financial Policy Committee

London School of Economics Systemic Risk Centre

30 October 2015

I am grateful to Omer Bugarinovic, Pavel Chichkanov, Rohan Churm, Kieran Dent, and Amar Radia for their assistance in preparing these remarks.

Last week, we at the Bank of England launched our new approach to stress testing the UK banking system.¹

We've published the approach, which has been evolving with our experience of stress testing, to give banks, markets and others as much clarity as possible.

And as I'll come on to explain, clarity about our approach has the potential to make stress testing an even more powerful tool to ensure banks have the strength to support the real economy in all weathers.

Our approach has three core elements - the three Cs:

It's **countercyclical**, and systematically so. In general, it will be more severe in what markets think are the good times and less so in the bad.

It's **consistent** with the capital framework for banks – an integral part of it, rather than an alternative.

And in assessing the results, it's **comprehensive** in approach – using a broad church of banks' own models, as well as our own.

What is stress testing and why do it?

But before I elaborate on those, allow me to take a step back and ask: what do we mean by stress testing and why are we doing it?

We mean assessing the impact on bank capital of a hypothetical bad economic scenario.

We want to know the damage the banking ship would sustain in an economic storm.

But we're not just interested in whether banks stay afloat. We want the system to be strong enough to continue to serve the real economy, even in the storm. In this year's stress test, that means expanding lending to the UK economy by 10% over the next five years, even in the face of a synchronised domestic and global slowdown.

Why is stress testing needed to ensure banks have the requisite strength? Aren't the baseline Basel III capital requirements sufficient to do the job?

¹ Available at <http://www.bankofengland.co.uk/financialstability/Documents/stresstesting/2015/approach.pdf>

After all, once fully implemented, the most systemically important banks could face capital requirements that are ten times tougher than they were before the crisis.² And these requirements are designed to reduce the probability of a systemic banking crisis to a very low level.

There is a reasonable debate to be had about whether a little more might be necessary given the costs of such crises.³ But with the advent of effective resolution regimes for failing banks and new requirements for loss absorbing capacity in resolution, which will dramatically reduce the economic consequences of bank failure, these requirements are not in the wrong ballpark.

However – and this is crucial – they are not designed for unusually risky banks in unusually risky situations. In particular, they are not designed for the more risky situations that emerge later in what we might call ‘the financial cycle’.

In the most risky situations – the apex of the cycle and often when markets think risk is at its lowest – banks would need to have capital ratios that are materially higher, to be confident of absorbing possible losses.

Requiring those sorts of capital ratios at all times - something not without its advocates – would certainly be prudent. But it would not make macroeconomic sense.

Costs of finance would be inefficiently high and the productive potential of the economy harmed.

So we need a capital regime that is both prudent and makes macroeconomic sense.

A capital regime that is macroprudential.

How? By changing capital requirements over time to match the risks on the horizon. When the environment is unusually risky, banks should hold more capital. But when it’s not, they shouldn’t.

Though it is far from the only one, stress testing can be a powerful tool for turning that principle into practice.

It tells us whether banks’ are taking ever bigger bets on the economy. And if the stress scenario can capture the risks banks face, the test can help to calibrate whether, and how much, additional capital is needed.

This is an increasingly important exercise. In the aftermath of a crisis, the financial cycle has been in a muted phase. But as it progresses into a more normal phase, we will have to devote much more attention to whether capital requirements should be increased to keep pace with the risks facing the system.

² See Caruana (2012).

³ A number of countries have exceeded minimum Basel III capital and/or leverage ratio requirements for systemically important banks; for example: Australia, Austria, Canada, Denmark, Finland, Hong Kong, Netherlands, Norway, Singapore, Sweden, Switzerland and United States.

Countercyclical

Which brings us to the first C of the new approach: countercyclical.

The stress scenario will become more severe as the risks get bigger and less severe as those risks either materialise or shrink.

This is ambitious and ground breaking. And if past cycles are anything to go on, we could be making the test more exacting just as lenders and their investors are thinking the world is less risky – that this time really is different. It will not be easy.

So how are we going to do it, first in principle, then in practice?

In principle we are making the design of the stress test scenario an exercise comparable to – but not the same as – economic forecasting.

Forecasting typically amounts to making a projection that is in the centre of the distribution of possible future outcomes – a projection that, as time unfolds, is as likely to be too optimistic as too pessimistic.

Producing a stress scenario – in our approach – is also making a projection. But it is not a forecast. It is a projection of a point in the tail of the distribution of possible future outcomes. In Figure 1, the forecast is point X; the tail projection – the stress scenario – is something like point Y.

It is a projection very likely to prove to be too pessimistic. The probability of it being too optimistic will be aligned with our appropriately low risk appetite for banks finding themselves so weakened that they crunch credit supply to the economy.

To produce that tail projection, we'll begin with the forecasts made by the Bank of England and IMF of the UK and global economies. By using distributions of past outcomes and analysis of forecast errors, we'll take the first step to making a projection of the point in the tail.

But it's the second step that's even more interesting, and even harder. We'll need to question whether the skew in the distribution of possible future outcomes is bigger than usual – more like profile B than profile A in Figure 2. If it is, the stress scenario – the same point Y in the tail of the distribution – should be more severe, as shown in Figure 2.

To gauge this, we will assess indicators of potential imbalances – in credit, in asset prices, and in household and corporate balance sheets. The emphasis is on potential because the exercise does not rest on spotting imbalances with any certainty. That is the strength of this approach.

When you're projecting a point in the tail, only a possibility that an imbalance might be present is needed to adjust it.

That's the principle. How might it have worked in practice if stress testing, and this approach to it, had been in place in the run-up to, and aftermath of, the financial crisis?

I'll use the example of the housing market, treating it in isolation for simplicity.

Figure 3 shows UK house prices relative to average earnings in that cycle. It also shows a range of estimates of the long-run level, drawing on rents, interest rates and past averages.

In the early part of the century, house prices were aligned with these metrics. There wasn't an indication of an imbalance. But even from such an apparently balanced position, history tells us that shocks happen. There is always a distribution of possible future outcomes.

The stress scenario – the projection of the point in the tail of the distribution – might have included a double-digit fall in house prices relative to earnings (the red arrow in the chart).

As the decade progressed, house prices continued to outpace earnings, and by more than could be explained by rents and interest rates. It was possible that an imbalance was emerging.

The probability of a correction was growing. The stress scenario should have got more severe. To calibrate how much more, we look at previous housing cycles and assess how their severity has related to the initial indicators of imbalance.

At what we now know to have been the peak of the cycle, the stress scenario would have included a fall in house prices, relative to earnings, of almost 30% (shown by the green arrow): twice as severe as earlier in the cycle.⁴

House prices then fell by almost 20%. And at that point, the severity of the test would have reverted back to where it started.

Now, projecting the middle of the distribution is hard enough. Projecting the tail is, as Sir Humphrey might say, "very brave". By their very nature, there isn't a large sample of tail events to go on.

But we've adopted this ambitious framework for two reasons.

⁴ This illustration treats the housing market in isolation. A real stress test would also include the knock-on effects of other economic imbalances. The 2014 stress test, for example, assumed that the UK current account closed and long term interest rates 'snapped back'. The knock-on effect of both of these elements of the stress scenario was incorporated in the scenario for the housing market, which included a fall in house prices of 35%.

First, the stress test is informing bank capital requirements, the setting of which has real economic consequences. It is incumbent on us not simply to ratchet up – or down – the severity of the stress scenario over time.

As policymakers, we have a responsibility to ensure that capital requirements change only because the risks change. As we move from stress test to stress test, the responsibility will be on us to explain the indicators that have caused us to change the severity of the scenario.

The side of the bed we exited on the morning of deciding the scenario should not feature in that dataset. It hasn't before. And this approach ensures it won't in future.

There is a second reason – more selfish on our part – for setting out a systematic countercyclical approach.

By looking at the economic situation, banks will increasingly be able to anticipate our actions. Although it can never be mechanical, our aim is for changes in the severity of the test to become more predictable.

With that, they can adjust their capital and business plans accordingly. They don't need to wait for the results of the test. They don't even need to wait for the test scenario to be announced.

Systematic policymaking can shape the behaviour of the banking system, hardwiring into its DNA a capital strength that counters the cycle.⁵

It is a big investment that will take many years to bear full fruit. But the prize is great. The development of systematic monetary policy in the past 20 years shows quite how great.⁶

That is not to say we are throwing away flexibility to stress test banks against risks that don't relate to economic imbalances. A central element of our new approach is the addition of a second stress scenario in every other year.

That 'exploratory' scenario is there to give us the opportunity to poke around in the corners, to test more structural risks to which UK banks might be exposed – risks that are not well captured by indicators of imbalances. Those might include structural changes in overseas policy regimes, the possibility of sustained deflation and low interest rates at the global level, or structural changes in particular industries.

⁵ By being more predictable, we can also help to ensure that, when more capital is needed in the system, it is added over a measured timeframe, reducing any costs of transition.

⁶ See Woodford (2003), for detailed discussion of this principle.

That biennial exploratory scenario ensures we retain sufficient discretion to keep up with all the risks we might want to test, while at the same time ensuring consistency through time in our annual scenario to capture the state of the cycle.

Consistent with the capital framework

The second C of our approach is that it's **consistent** with the capital framework.

Our approach to stress testing is not about building an alternative framework of capital regulations. We are integrating stress testing into the framework that's already in place.

Here's how.

Capital requirements fall into two sets: minimum requirements that must be held in all circumstance, however stressed, and buffers on top of that minimum that can be depleted to absorb and cushion the impact of stress. That 'capital stack' is shown in Figure 4.

Those buffers really are there to be used. It would be self-defeating to require them to be maintained at all times or to be rebuilt at breakneck speed. Their usability is vital to avoiding the crunching of credit supply to the real economy.

Stress testing informs how big that buffer needs to be to deal with future storms. As risks grow and storm clouds gather, the stress scenario will get more severe and buffers will increase. If the storm hits, the buffers will be used.

Matters are complicated by the fact that the total capital buffer for each bank has many components.

All banks have a capital conservation buffer.

That can be augmented across all banks by macroprudential authorities (in the UK, the Bank of England Financial Policy Committee).

And it can be augmented further on a bespoke basis by microprudential supervisory authorities (the Prudential Regulation Authority of the Bank of England).

In the UK, we're fortunate. All the authorities sit in one place – the central bank – so the co-ordination to avoid over and underlap is quite achievable. But for the Kremlinologists, here's how we'll do it.

First, the macroprudential authority will identify in the bank-by-bank results of the test the common-system wide impact.⁷

Where that is greater than the conservation buffer and any supplementary macroprudential buffers in place, it will inform a decision on whether, and how far, to raise those macroprudential buffers, including the countercyclical capital buffer.

Once these system-wide components are set, the microprudential authority will, using its supervisory judgement, set bank-by-bank top up buffers.

Such a top-up will be appropriate for banks whose balance sheets are more sensitive to the stress than others, such as those who use highly pro-cyclical risk-weight models.⁸

All of this will be going on inside the Bank of England. To the outside, we will be one central bank, working in concert to set a total capital buffer for each bank that's big enough to deal with a storm if it hits.

Our new approach is also consistent with the soon-to-be-implemented additional capital buffer that globally systemic banks must hold. Four of the seven banks currently in our stress test will be subject to an additional buffer, but each to a different size.⁹

The principle behind these buffers is to ensure that, the more systemic is a bank, the bigger the stress it can absorb. The tolerance of risk of failure is lower for systemic banks.¹⁰

We had two choices to align stress testing with this principle. The first was to subject each systemic bank to a different severity of stress scenario – a projection of a different point in the tail of possible outcomes.

The second was to subject banks to a common scenario reflecting our risk tolerance for less systemic banks, but to hold globally systemic banks to a higher standard by not allowing the systemic buffers to be used to deal with the stress in our test.

For practical reasons, we chose the second.

But this does not mean we think the additional capital buffers that systemic banks must hold are unusable in a real stress. Quite the opposite: like the other elements of the capital buffer, they are there to be depleted if necessary.

⁷ This exercise will draw on the full set of results of the stress test, including results from small and medium sized banks running the scenario as part of the Internal Capital Adequacy Assessment Process (ICAAP).

⁸ An issue highlighted in the 2014 stress test results, available at <http://www.bankofengland.co.uk/financialstability/Documents/fpc/results161214.pdf>.

⁹ See FSB (2014).

¹⁰ Cunliffe (2014) sets out the challenges involved with the failure of systemically important institutions.

It simply reflects that we want to have confidence that these banks have a buffer big enough to withstand a stress that's more severe than the one we are applying to all banks.¹¹

And by doing so, we have a stress testing approach that's completely consistent with the principles of the capital framework for globally systemic banks.

Comprehensive in modelling approach

The third C of our new approach is that it is **comprehensive** in its approach to modelling the impact of the stress.

We'll continue to ask banks to perform the test themselves, on their own balance sheet. We are not intending to apply our own models exclusively.

This allows banks' own stress testing capability to be put to the test. We want assurance that they can test themselves, for the specific risks they face, and do so regularly.

It also ensures a wide range of models is brought to bear on each test. Each bank brings its own approach, so right now, that's seven modelling approaches in the test. We can compare them, pick the best for the problem at hand, and adjust the results of the others.

Our focus will be on developing models in areas that all banks find most challenging. The modelling of net interest income is one such area. But we are not seeking to check every part of banks' own efforts every year. Our resource is better spent on capturing what we can't expect the banks to capture: the system-wide dynamics of the stress.

These are amplifiers and feedbacks that were so evident in the financial crisis.

The mechanisms that turned so-called six-sigma events into day-to-day events.

The mechanisms that prompted Ben Bernanke to argue in 2008 that central banks needed to "widen our field of vision".¹²

I'd highlight three mechanisms that should be brought more fully into our field of vision.¹³

¹¹ The supervisory response to a breach of Pillar 1 and Pillar 2A CET1 requirements in the stress will be more intensive relative to a failure to meet systemic buffers. For banks that fall below their minimum Pillar 1 and Pillar 2A CET1 capital requirement in the stress, there will be a strong presumption of an intensive supervisory response to rebuild capital. Systemically important banks that fall into their systemic buffers, but not their minimum CET1 capital requirements, will still be expected to strengthen their capital positions. But the supervisory response will be less intensive across one or more dimensions, including the size, nature and timing of required remedial actions (see Section 1.4 of The Bank of England's approach to stress testing the UK banking system).

¹² See Bernanke (2008).

First, the feedback loop between weakened capital and higher funding costs, not just for individual banks, but across banks with similar business models.¹⁴

Second, the feedback mechanisms associated with banks' liquidity positions.

Like capital buffers, liquid asset buffers are there to be used in stress. And central bank facilities are open for business.

But in a stress, weaker banks may see the tenor of their funding shorten, there may be a flight to (perceived) quality by investors, and derivative counterparties can call for more collateral. All of this will place demands on liquid asset buffers.

Those buffers may not be sufficient to absorb the full impact. Where that's the case, we'll need to factor in the impact on capital strength of asset firesales and, where appropriate, the impact on funding costs of use of central bank facilities.

Third, our stress testing field of vision should begin to extend to non-bank parts of the financial system.

Investment funds are one such, but far from the only, part.

They have grown rapidly in importance since the crisis and in many cases are offering – or are perceived to offer – the ability to redeem shares at short notice.

On the back of that, the Bank of England and the FCA have together been surveying these funds and assessing the risks associated with their growing importance.

Individual funds rightly manage their liquidity and test their ability to meet redemptions under stress. They tend to have contingency plans. They have little debt. And their investors bear the market risk. So there's not really a question of whether they can withstand a storm.

But a macroprudential stress test would ask a different question.

It would ask: will investors and fund managers behave in a similar way after a correction in financial markets?

Will investors try to use the redemption possibility they're being offered *en masse*?

¹³ For a full review of these mechanisms and the approaches to deal with them see Demekas (2015).

¹⁴ In this, we have made progress already, but more is needed to put this into practice see Aikman (2009).

And if that means funds together try to sell into the same markets at the same time, how will that magnify volatility, disrupt the real economy, and feed back to the core of the financial system?

So their inclusion in any stress testing would be very different from the way we treat banks. But if the investment fund community is collectively engaged in activity that could create systemic risk – and that is a matter for further analysis at domestic and international level – then those activities should be within our field of vision.¹⁵

Conclusion

So there is a great deal of work to develop the approach in future.

But our countercyclical, consistent and comprehensive plan will see the UK's bank stress testing regime fit for the next stage of the cycle.

It's an approach that increases the predictability of our actions; that hardwires stress testing into the DNA of the banking system.

An approach that slots stress testing into the capital framework, not one that rivals it.

And above all, it's an approach that gives the economy what it deserves:

A banking system that's prudent, makes macroeconomic sense, and that supports it in both calm seas and severe storms.

Thank you.

¹⁵ The Financial Stability Board, at its meeting in London last month encouraged funds to “assess their ability individually and collectively to meet redemptions under difficult market liquidity conditions”, see FSB (2015).

References

- Aikman, D, et al, (2009)**, “Funding liquidity risk in a quantitative model of systemic stability”, Bank of England Working Paper No. 372, available at:
<http://www.bankofengland.co.uk/research/Documents/workingpapers/2009/wp372.pdf>
- Bernanke, B, (2008)**, “Reducing Systemic Risk,” remarks at the Federal Reserve Bank of Kansas City’s Annual Economic Symposium, Jackson Hole, WY, available at:
<http://www.federalreserve.gov/newsevents/speech/bernanke20080822a.htm>
- Caruana, J, (2012)**, “Building a resilient financial system”, available at:
<http://www.bis.org/speeches/sp120208.pdf>
- Cunliffe, J, (2014)**, “Ending Too Big to Fail – progress to date and remaining issues”, available at
<http://www.bankofengland.co.uk/publications/Documents/speeches/2014/speech727.pdf>
- Demekas, D, (2015)**, “Designing Effective Macroprudential Stress Tests: Progress So Far and the Way Forward”, IMF Working Paper WP/15/146, available at:
<http://www.imf.org/external/pubs/cat/longres.aspx?sk=43043>
- FSB (2014)**, “2014 update of list of global systemically important banks (G-SIBs)”, available at:
http://www.financialstabilityboard.org/wp-content/uploads/r_141106b.pdf
- FSB (2015)**, “Meeting of the Financial Stability Board in London on 25 September”, available at:
<http://www.financialstabilityboard.org/wp-content/uploads/September-Plenary-press-release.pdf>
- Woodford M, (2003)**, “Interest and prices”, Princeton University Press, available at:
<http://press.princeton.edu/titles/7603.html>

Appendix

Figure 1 : Stylised conditional distributions of expected movements in a variable




Figure 2: Stylised conditional distributions of expected movements in a variable


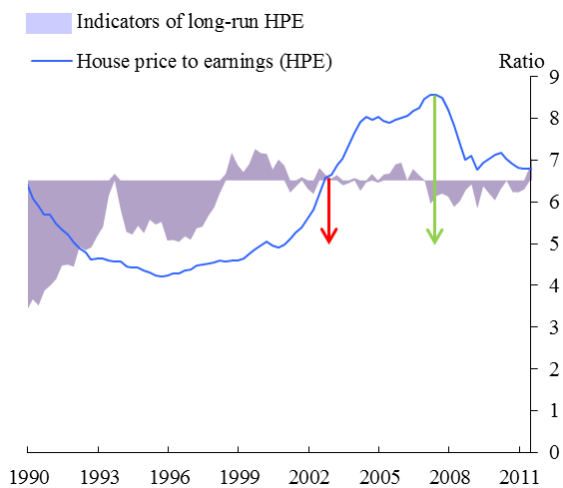


Figure 3: UK house price to earnings ratio



Sources: Bank of England, Halifax, Nationwide, ONS and Bank calculations.

Notes: House price to earnings ratio calculated as the ratio of the average nominal house price and average annual earnings. Indicators of long-run HPE include post-1990 median HPE and HPE estimated using a dividend discount model (DDM) that takes account of levels of rents and interest rates.

Figure 4: Components of the regulatory capital framework

