Bufferati - speech by Sam Woods

Given at City Week 2022

Published on 26 April 2022

Sam Woods looks at the rules that are in place to make sure banks have enough capital put aside to withstand economic shocks. He sets out a radical, alternative approach that he argues would be much simpler.

Speech

Concept cars

I have always wondered why car manufacturers produce concept cars – those futuristic prototype vehicles, not yet ready for production but good for getting the manufacturer's brand noticed. Is it to open the minds of their customers to new ideas coming down the track? Is it to create a halo effect, so that a middle-aged parent driving their family saloon can imagine they are at the helm of a sleek, spaceship-like craft from the future? Or is it just that it's really fun for the car companies' engineers to invent these concept cars, and for board members to check them out during the coffee break from their interminable audit committee meetings?

Well, perhaps a car company executive will read this speech and let me know. I concede that this is very unlikely, and indeed I would be very reluctant to buy a car from any company whose executives spend time reading my speeches! So I'm going to guess that it's a combination of all the above – but for my purposes today my main focus is the one about having fun.

Because what could be more fun than financial regulation? Or more specifically, what could be more fun than examining the way we currently set capital requirements for banks? I realise that the true answer to this question is a staggeringly long list of more entertaining things you can spend your time doing, including checking out the latest concept cars on YouTube if you are into that sort of thing. So in the interests of time I will move on swiftly.

Buffer usability

As many of you will know, there is a lively debate going on amongst regulators about what we call "buffer usability" – the extent to which banks are in reality willing and able to use their buffers of capital to keep lending to the real economy in a downturn. Learned colleagues on the Basel Committee, including Vicky Saporta and David Bailey from my own team, have put the current model of capital requirements back onto the assembly line and are painstakingly examining each section and considering how it has performed as part of an evidence-based evaluation of post-crisis reforms.[1] They are also giving a little thought to the complexity of the current design, which I can tell you from having driven it a lot is very considerable.

This is good work, very ably presided over by the Chair of the Committee Pablo Hernández de Cos, and I fully support it. There is also a recognition that unnecessary change is undesirable and the bar for evidence is therefore high, given we have just put the finishing touches on the post-crisis reforms – reforms which the PRA is fully committed to implementing faithfully. But I have found myself thinking it might be a good idea to introduce, on the next door assembly line as it were, a new concept car version of the capital stack – radically simpler, radically usable, and a million miles away from the current debate but which might prove instructive over the longer term. At the very least this might amuse my colleagues on the Basel Committee as they toil away on the serious business of evaluating the current model, perhaps in a coffee break away from the mystical intricacies of the CCoB, the CCyB, GSIB buffers, DSIB, O-SII, ring-fence and PRA Buffers (depending on which part of the world you come from), MDA restrictions, non-buffer elements of Pillar 2 and of course Pillar 1.[2]

Now of course by far the most important thing about a concept car is that it must look either incredibly cool or incredibly bizarre, and preferably both. The graphics team in the Bank of England has many talents but I concede defeat on this aspect, if we are up against things like the Bentley EXP 100 GT or the Mercedes Vision EQ Silver Arrow. But the next most important thing is to have a good name, and as this is all about buffer usability let's call our concept capital stack the "Basel Bufferati". As you can see, we do not employ a marketing department! I should point out though that I have no interest whatsoever in designing a weak car – like my colleagues on the Prudential Regulation Committee (PRC) and Financial Policy Committee (FPC) I am fully committed to maintaining robust prudential standards.

Design features

While the capital regime is fiendishly complex, its underlying economic goals are fairly simple: ensure that the banking sector has enough capital to absorb losses, preserve financial stability and support the economy through stresses. In developing the Bufferati, my guiding principle has been: any element of the framework that isn't actually necessary to achieve those underlying goals should be removed. The Bufferati is as simple as possible, but no simpler.

With that mind, my simple framework revolves around a single, releasable buffer of common equity, sitting above a low minimum requirement. This would be radically different from the current regime: no Pillar 2 buffers; no CCoBs, CCyBs, O-SII buffer and G-SiB buffers; no more AT1.[3]

Let's get under the hood of the Bufferati. There are a few key elements – some of which closely resemble the current regime, and others which are quite different.

At the core of this concept is a single capital buffer, calibrated to reflect both microprudential and macroprudential risks and replacing the entirety of the current set of buffers.[4] After the global financial crisis, we rightly opened our eyes to the virtues of macroprudential policy: pre-crisis regulation had neglected system-wide risks. We now have a robust system of micro- and

macroprudential regulation, where macroprudential policy takes a bird's-eye view of the system as a whole and microprudential gets deep into the specific risks to the safety and soundness of firms.

You can of course have circumstances where institutions charged with these responsibilities conflict, and indeed I have experienced this myself. The hypothetical example most often cited is the risk of microprudential authorities not allowing banks to use their buffers in a stress, which might be good for the resilience of those individual institutions but would have an adverse impact on the wider system. But I would consider this a failure of microprudential policy rather than a fundamental conflict between micro and macro objectives, because in my view not allowing buffers to be used in a downturn will be bad for individual firms' resilience as well as the system as a whole. If banks have to preserve their capital ratios during a stress, they may do so by fire-selling assets and cutting back on lending to the real economy. In aggregate, this is the kind of behaviour that turns recessions into depressions, and it would be myopic for bank supervisors to welcome such an outcome.

So while individual macro and micro policymakers may disagree on particular issues, I think it's possible to make too much of the distinction between micro- and macroprudential regulation. They both have the same fundamental goal: a financial system that serves the real economy, including through stresses. They aim to avoid banking crises, disorderly failures or adverse credit supply shocks that can amplify downturns and turn recessions into depressions.

It's therefore worth asking – do we really need a raft of separately calibrated 'macro' and 'micro' instruments to achieve these goals? For the Bufferati, I envisage each firm's capital requirement as a single number, which would be informed by both macro and micro analysis:

- Macroeconomic cost-benefit analysis would be used to set the overall level of capital requirements in the system as a whole.[5] That number would move up and down as systemic risks build and subside.
- Microprudential analysis of firm-specific risks would set the specific buffer for each individual firm – but the average of those buffers would need to stay within the macro-economically optimal range.[6]

As an illustration – if you ran a macroeconomic analysis that said the average capital level should be 14%, then the supervisor might set the capital requirement for a firm with riskier exposures at 16%, while another lower-risk firm might get 12%, and so on. I think this is a pretty clean division of responsibilities, and one that reconciles macro and micro policy more effectively than giving the two authorities separate buffers to play with.[7] You would need governance arrangements in the case that the authorities disagreed amongst themselves about capital levels, such as the directions power the FPC has in the UK, but this strikes me as a more efficient arrangement than having multiple buffers serving different purposes and masters.

The next key feature of the Bufferati is replacing all thresholds, triggers and cliff-edges with a judgement-based 'ladder of intervention'. We've learnt over time that automatic thresholds for supervisory intervention come at a real cost – for example, firms will do what it takes to stay above their MDA trigger, including cutting lending to the real economy which rather defeats a main part of the purpose of capital regulation.

In our new world, there would be no automatic consequences to dipping into the Bufferati, but if that happened the firm would be expected to have a plan to rebuild their capital resources. What needs to be in that plan would vary widely depending on the firm-specific and/or macroeconomic circumstances which had led to the firm entering its buffer, including whether some of the buffer had already been released by the authorities. Both micro and macro authorities would have a range of tools which could be deployed as needed in a ladder of intervention, but these tools would never be triggered automatically by hitting a pre-determined threshold. For instance, in an extreme and sudden systemic shock such as the arrival of Covid-19 in Europe and the US in early 2020 the authorities might choose (as they did) to restrict dividends at a system-wide level, but in a situation where one firm experienced a large but idiosyncratic trading shock the microprudential authority would take actions tailored to the safety and soundness of that firm.[8]

Now, I say replace "all" thresholds in the capital stack but the Bufferati engineers have been debating whether to keep just one: a minimum requirement sat underneath the buffer, to serve as an airbag to protect what really matters during a crash. The most radical engineers think we should do away with the airbags too, but their more cautious colleagues have less faith in the new technology and think we need to keep this basic safety feature.

Now, there is nothing magical about a minimum requirement. But it does put the supervisor in a stronger position, from a legal point of view, to take the most severe mitigating actions as needed – including to preserve capital for resolution if the collision looks unavoidable.[9] But (at least in a UK context) nothing automatic happens when a bank breaches its minimum requirement – indeed, I have direct personal experience of supervising two banks which ran for some time well below their minimum requirement, but were successfully recapitalised and are now alive and well. This was of course a very uncomfortable ride in both cases but they were rides worth taking.

I conclude that the Bufferati would benefit from a simple minimum capital requirement underneath the buffer, as a basic safety feature with which to strengthen the supervisor's hand in extreme cases. But this minimum should be set at a low level, to leave maximum space for the buffer to do its job, and breaches of it should not have automatic consequences.[10]

This leads neatly to the Bufferati's third key feature: like the CCyB in the current framework, the entire buffer is potentially releasable in a stress. As regulators, we want banks to build up capital in good times, so that banks are in a strong position to deploy capital to support lending in a downturn. Releasing capital requirements in a downturn allows banks to do this more easily and readily, as it means they don't need to worry about entering buffers and triggering restrictions or

other supervisory actions.

Under the Bufferati, when systemic risks crystallise, the authorities could cut the system-wide average capital requirement to whatever level they deem necessary to avoid an unwarranted contraction in credit conditions and give firms room to manoeuvre as they go through the stress. A key benefit of this simpler framework is that we would be able to make very substantial cuts, far beyond the cuts to the CCyB rate that the UK's Financial Policy Committee and some other countries' macroprudential authorities have made in recent episodes.[11] This would have real benefits: early evidence from the Covid crisis suggests that releasing buffers is an effective way to avoid costly deleveraging by banks.[12] This also points to the importance of building up strong buffers well in advance of the shock hitting.

Of course, even if buffers were cut to zero, there is a risk that banks will feel under pressure from the market to preserve their capital ratios in a stress. Below a certain level of capital – call it the 'market imposed capital requirement' – banks will struggle to fund themselves. But in cases where the regulatory limit is above the market-imposed requirement, releasing the buffer will give banks more headroom. And releasing regulatory requirements can also lead to an easing of market-imposed requirements, for example by giving markets assurance that the banks will not face regulatory consequences for allowing their capital levels to fall.

The fourth key feature of the Bufferati is that all requirements are met by common equity. Common equity is the quintessential loss-absorbing instrument and is easy to understand. Instruments like AT1 and 'contingent convertible' debt have their place in the current framework but they introduce complexity, uncertainty and additional 'trigger points' in a stress and so have no place in our stripped-down concept vehicle.

The fifth key feature of the framework is a mix of risk-weighted and leveraged-based requirements. In the spirit of the Bufferati, one should ask whether we need two ways of looking at capital – for instance should we move to a purely leverage-based system, as some critics of the current regime have advocated? That would certainly be a simpler regime. But even in this high-concept vehicle, I don't think we can live without risk-sensitivity in capital requirements. Building a capital framework that doesn't allocate more loss absorbency to unsecured, sub-prime or otherwise riskier assets would be highly inefficient – we would probably need to run the whole system at a materially higher capital level to get comfortable with this. And such a framework would incentivise firms to take bigger risks to maximise return on equity. At the same time, risk-weights can clearly be badly wrong and excessive leverage is a danger in and of itself regardless of asset quality, and I therefore conclude that a leverage-based measure will also always be needed.[13] I envisage that the structure of the Bufferati leverage requirement would look similar to the risk-based one: a low minimum with a single releasable buffer.

Crash tests

Last but certainly not least, stress testing would be the central analytical input that ties the regime together. Conceptually, you can see stress testing as an expression of our risk appetite: it tells you how severe a stress an optimally-capitalised banking system is expected to weather. And it puts individual portfolios under the microscope to help us identify which firms are carrying more severe risks.

Our current approach to stress testing large banks in the UK revolves around a single 'annual cyclical scenario' (ACS) stress test, where we test banks' ability to keep their capital levels above a pre-set hurdle rate in a severe but plausible economic downturn. The stress test is a large exercise for ourselves and for firms, and is one of the most richly valuable pieces of the regulatory framework. Similar approaches, sometimes with a small handful of scenarios instead of just one, are in place in other jurisdictions.

But in the spirit of the Bufferati, I wonder if another world is possible – one where stress tests are just as robust, but simpler to run, more frequent, and cover a much wider range of economic outcomes. We could move away from a single annual scenario in favour of thinking holistically about our risk appetite across a range of scenarios. After all, in the European insurance framework firms run literally thousands of scenarios to test their resilience – are we so sure that one or two every year or so is the best answer for banking? One can imagine a world where we set a hundred scenarios, based on the distribution of possible shocks to the sector, and set a risk appetite across them. Rather than just pick a point on the distribution of scenarios, as we do in Europe for insurance companies, the overall level of a capital would be set by the macroprudential assessment above – but stress testing would let us know how much of the distribution of scenarios banks can withstand. This would be a useful cross-check on the top-down macroprudential judgement on capital levels, would identify outliers and could also be one way of allowing the authorities to adopt a higher risk appetite for smaller banks by requiring systemic banks to be robust to a wider set of scenarios.[14]

In developing this enhanced buffer framework, it is possible that a lot of the sophistication which currently resides in modelling risk-weights would move into stress testing. Indeed, this is already happening in Europe because the move to expected credit loss accounting has necessitated the development of more IRB[15] -like modelling techniques to inform the impairment provisions which feed into the stress test. Ultimately, in the spirit of simplification one could perhaps argue that with this migration one could move the system to standardised risk-weights, and rely on stress testing to deliver the sophistication which currently comes from modelling risk-weights – but as this question is not central to the Bufferati's design I will leave it to the next generation of engineers to settle.

Unmoored from reality?

So in summary, the Bufferati has:

1. A single capital buffer, calibrated to reflect both microprudential and macroprudential risks.
2. A low minimum capital requirement, to maximise the size of the buffer.
3. A 'ladder of intervention' based on judgement for firms who enter their buffer – no mechanical triggers and thresholds.
4. The entire buffer potentially releasable in a stress.
5. All requirements met with common equity.
6. A mix of risk-weighted and leverage-based requirements.
7. Stress testing at the centre of how we set capital levels.

Clearly this framework diverges heavily from current international standards. But it's worth reiterating that any capital framework – complex or simple – will only be successful if it is based on strong, common international standards. There is no reason why a global standard needs to be complex or overladen with buffers and thresholds. One can imagine a global Bufferati designed and overseen by the Basel Committee, with international agreement on the design features listed above (including the overall capital level and risk appetite, the approach to risk-weights and a common framework for stress testing) and national authorities responsible for implementation, including by setting firm-specific requirements and judging the overall level of systemic risk in the domestic financial system.

To maintain a level playing field globally, these judgements would need to be made in a clear international framework including a macroeconomic anchor point for overall capital levels in normal times, guidance on things like what constitutes a stress, approaches to stress testing and speed of buffer rebuild, and how much loss-absorbing capacity is needed in resolution. In a nutshell, to oversee the Bufferati we would need an international framework that is clear and consistent, but also judgement-based. This relies on our ability to agree a common set of objectives and approach to analysis, so that individual jurisdictions' judgements are made in a common framework. It also relies, inescapably, in a certain amount of international trust and goodwill, that we can achieve this without a system of automatic triggers and thresholds – otherwise we would either end up with fragmentation, or re-introducing all the complexity of the old regime.

Now of course concept cars often come to nothing, maybe partly because the really good ones are totally unmoored from reality. I suspect that expert recipients of this speech will conclude that the Bufferati is just such a vehicle, miles away from reality and not even very nice to look at. How could we run the system well without a dedicated explicitly macroprudential buffer? How can we ensure safety and soundness and market confidence without another buffer specifically for conserving capital? How could it be right not to have a special extra buffer for larger banks, given the obviously greater potential they have than smaller banks to inflict damage on the economy? Clearly this particular car's designers have lost their minds. Surely we need a steering wheel, an engine, a set of wheels?

Well, they do have a point of course, and I guess this particular assembly line will be quickly repurposed to producing a mass market electric estate car. But as I prepare to rebuild it I take a quick look at some of the other assembly lines already chugging away in the financial regulatory factory. Two in particular catch my eye.

First, the liquidity coverage ratio (LCR) assembly line. Now this is interesting, despite the boring name.[16] This vehicle is one giant buffer. No “LCoB”, no “GSILB”, no “LCYB”, no minimum. My experts will point out quite rightly that capital is not the same as liquidity, and that the theology behind the development of the LCR is quite different. Alternatively, one might argue that the LCR is the best that could be agreed, but in an ideal world, it would have some of the features that we have in the capital stack. Most obviously, one could point out that the buffer usability issues in relation to the LCR seem every bit as bad, indeed perhaps worse, than those we have on the capital side.[17]

Still, I find it striking that we have such a radically simpler, sleek vehicle already rolling off the adjacent assembly line.[18]

Second, on the other side of the shop floor are a mysterious and eminent group of engineers, dressed as druids. These are the actuaries and in Europe they have produced an incredibly cool and popular car known as Solvency II. Now, I can see that I may not be onto a winning argument here! But it is interesting to me that we have another set of core financial institutions whose capital requirements are also, like the LCR, set in the form of one large buffer – plus a low minimum. Of course, banks are not insurers and insurers are not banks, not least from a macroprudential perspective. And there are plenty of issues with this vehicle – not least the dizzying complexity of the stress testing models used to set the level of the buffer for many firms.

But as with the LCR, I am intrigued that we have such a thing in an adjacent part of our world, and wonder if perhaps it suggests that the Bufferati is not quite so unmoored from reality as it may first appear.

On the road

With that, I must leave the regulatory factory and get back on the road.

I settle into my familiar, hard-wearing Basel 3.1. I have no desire whatsoever to weaken the vehicle, and I had to save up for a decade after the financial crisis to buy my Basel 3.1 so I am fond of it. I didn’t like the convertible model but the sun is shining so I open the sunroof. I look at the dashboard, consult the user manual to remind myself which of the numerous dials is the fuel gauge, and pull into the next petrol station to fill up.

Next to me at the electric charge point is a long, sleek vehicle seemingly moulded from single sheet of titanium and with a set of drone rotors in each corner. An immaculately styled regulator comes out of the shop with an ice-cold latte in hand. She hops into the driver’s seat, powers up

and takes off into the future in a brand-new Bufferati.

I am grateful to Hugh Burns, Vicky Saporta and many other colleagues for their help in preparing this speech.

---

1. See Vicky Saporta's June 2021 **speech on Emerging prudential lessons from the Covid stress**
2. For those less familiar with these acronyms: the capital conservation buffer (CCoB), countercyclical capital buffer (CCyB), global systemically important bank (GSIB), domestic systemically important bank (DSIB), other systemically important institutions (O-SII), and maximum distributable amount (MDA).
3. Alternative tier 1.
4. The buffer would be fully disclosable.
5. There have been various previous attempts to quantify the optimal level of system-wide capital requirements, for example the Basel Committee on Banking Supervision's 2010 **assessment** of the long-term economic impact of stronger capital and liquidity requirements, or a 2015 Bank of England **Paper on 'Measuring the macroeconomic costs and benefits of higher UK bank capital requirements'**.
6. In practice, some firm-specific risks might not be captured in macroeconomic analysis – for example failings in governance and controls. For this reason, it would be important to set the system-wide level as a range, so that microprudential judgements are not unduly constrained.
7. We would also want to take account of how costs of failure can differ across firms. But rather than needing a separate 'systemic buffer', we could simply give systemic firms a higher average capital requirement as illustrated in the stress testing section below.
8. Of course, there is a risk that even without automatic triggers and thresholds, firms will want to stay out of their buffer because of the feared market reaction from dipping into it. That's why releasability is such a key feature of this design – as discussed further on.
9. For the purposes of this speech I am focusing on going concern capital requirements, though any framework will also need to ensure banks have sufficient loss absorbing capacity to be resolvable. In the Bufferati simplified bail-in debt requirements could either be driven by minimum requirements or, if that appeared too low, a simple % (such as 50%) of the entire going concern capital stack.
10. If applied to the current UK framework, it would mean doing away with the current distinction between Pillar 1 and Pillar 2A minima, where the latter covers firm-specific risks. We would need to ensure that the risks currently captured in Pillar 2A are addressed elsewhere – for example in the calibration of the buffer, in risk weights or in stress tests.
11. In judging how far to cut the buffer in response to a given shock, the authorities would weigh up the benefits of supporting the economy in the near term, versus the temporary reduction in resilience against further shocks.
12. This evidence is set out in Vicky Saporta's 2021 speech. In brief, firms that were more exposed the FPC's 2020 CCyB cut increased their CET1 ratios by less during the pandemic. This analysis was carried out by a team of Bank staff including Aakriti Mathur, Matthew Naylor, Aniruddha Rajan, Katie Low, and Damien Lynch.
13. 'Model risk' is most associated with banks' own internal models, but even standardised risk-weights set by regulators can be undercooked – an obvious example is the zero risk weight assigned to Greek sovereign debt before the financial crisis.
14. This implies that we would be publishing scenarios in which we expect that some banks would fail. This would be a change but in any framework without 100% capital requirements there are circumstances where firms can fail – something that can get lost when you centre your regulatory approach around a single flagship stress test.
15. Internal Ratings-Based.
16. Apologies to British carmaker TVR, whose cars do look very nice despite the acronym.
17. This may in part reflect the fact that there is no minimum requirement for the LCR.
18. For more on this, see the PRA's recent discussion paper **DP1/22 – The prudential liquidity framework: Supporting liquid asset usability**.