BANK OF ENGLAND

Speech

Credit union meets robot

Speech given by
Sam Woods, Deputy Governor for Prudential Regulation and Chief Executive Officer, Prudential Regulation Authority

Mansion House, London
Thursday 24 October 2019

I am very grateful to Hannah Schraer for her help in preparing this speech.

Introduction

Earlier this year I gave a speech, which fortunately for most of you in this room you did not have to sit through, on competition. I avoided the temptation to list all of the things we've done to facilitate competition and instead posed a set of seven questions which I think are more awkward for us.

The three-line version of that speech is this: we've done a lot to lower prudential barriers to entry into the banking sector, with the result that we've got a lot of new banks. But have we done enough to lower the equivalent barriers to growth, given no little bank has recently become a really big bank?

Since then I have unleashed the competition team within the PRA to roam about the institution digging up whatever they can find on the barriers to growth issue. They have already come up with two things. First, with the help of a PRA robot they have for the first time been able to map all the barriers, and show how complex our rules are.[1] And second, together with our credit union supervisors they have taken a fresh look at capital requirements for that sector and proposed some quite radical changes in order to promote the growth of successful firms.[2]

Barriers

The first thing we've done is take a fresh look at the various thresholds which arise in our regulations. To date this debate has been massively dominated by two features of the regime: the ability of firms to model their capital requirements (using IRB) rather than use a standardised approach, and the need for firms to issue bail-in debt (MREL) once they reach a certain size.

These are both very important issues. On credit risk, we have taken three very important steps in recent years: the introduction of our modular approach for IRB applications; the introduction of our "holistic" approach to P2A-setting for eligible firms on standardised; and our success in getting Basel to agree to lower the Standardised Approach risk-weight for low-LTV mortgages very substantially, from 35% to 20% for firms on standardised. MREL, set by the Bank as Resolution Authority, is set to ensure that a bank in resolution can meet going concern regulatory capital standards, so any change in regulatory capital standards will automatically be reflected in MREL. We continue to take a close interest in both of these areas, but they have crowded out discussion of everything else in the barriers to growth area – so today I want to fill in the rest of the picture.

Now remember: these regulatory thresholds are a direct result of decisions to exercise proportionality by exempting smaller firms from bits of regulation. So they are a good thing from a competition point of view.

[1] The robot's research, including code, is being published today. See: Amadxarif Z., Brookes J., Garbarino N., Patel R. and Walczak E. (2019), 'The Language of Rules: Textual Complexity in Banking Reforms', Bank of England Staff Working Paper.
[2] My focus today is on deposit-takers, but there are also of course important competition questions for us to examine on the insurance side.

They also ensure that the regulation of different-sized firms is commensurate to the risks that they pose to the financial system, and are therefore sensible from a risk perspective as well as a competition one. But thresholds do create the "slope" issue: as firms grow, they come up against these thresholds and have to clamber over them. The key question for us is: does the design of these thresholds, and the way they come together across different bits of regulation, present unnecessary barriers to firms looking to climb the mountain? Have we created any unnecessarily perilous cliff-edges?

Now, a few years ago the way we'd have investigated this would have been to get an analyst from our Policy Directorate to enter the mighty forest of UK financial regulation and hunt down any and all thresholds for our examination.

However, last year – quietly and without great fanfare – our team started to turn this forest into a dataset better suited for textual analysis. I like to call this a machine-readable Rulebook but this upsets the experts in the PRA, so instead you should think of it as us having chopped up all of the rules into bite-sized elements and put each of those into a cell in a giant database. Our team is now publishing the code used for this effort, and a Staff Working Paper with early research findings from the data. One great benefit of this is that rather than tasking an unfortunate analyst with the Herculean endeavour of mapping all the thresholds, we are able to use a PRA robot to do the job instead.[3]

The PRA competition robot has returned from its daunting mission with interesting findings.

First, there are 53 thresholds for deposit-takers, spread across all the main policy areas but with a concentration in capital (17), reporting (11) and remuneration (8).[4]

Second, 13 unique metrics are used to set thresholds. The most common, unsurprisingly, is simply assets – 26 of the 53. But even within these asset-based ones 16 different values are used to set the actual threshold.

Third, 20 of the thresholds result from EU regulations and/or Basel definitions. The other 33 are things we have done ourselves here in the UK.

Fourth, if we crudely translate non-asset thresholds into an asset equivalent in order to be able to map all thresholds on one scale, then most are concentrated in the area below a size corresponding to £50b in assets or deposits.

[3] This bears out the prediction made last year by former PRA Executive Director James Proudman that the PRA would soon develop a form of 'cyborg supervision' involving humans and machines working closely together.
[4] To identify thresholds, the robot ran a series of textual searches on the dataset. The search combined quantitative values (£, €), expressions related to thresholds (eg. less/more than, greater/smaller than, up to); and typical metrics used for thresholds (eg. assets, balance sheet, deposits, accounts). It also included thresholds, such those on executive pay and exposures, that are only indirectly linked with size. The results were then cross-checked and cleaned by subject matter specialists.

One could make various observations about these findings. But one obvious question to ask is whether or not this is too complicated a picture. While it is a good thing for small firms that we exempt them from various bits of regulation, is there a case for making the gradient a bit easier to discern and climb as they grow? Should we rationalise and harmonise any of these thresholds? We are now looking at that question and have an open mind on it.

But in a way the more interesting thing the robot brought back to us was not these findings, but simply a clearer picture of what the gradient is for small firms which we are therefore publishing today for the first time. Here it is:

Chart 1a: Thresholds in UK prudential regulation, translated onto an assets/deposits basis (deposit-takers)

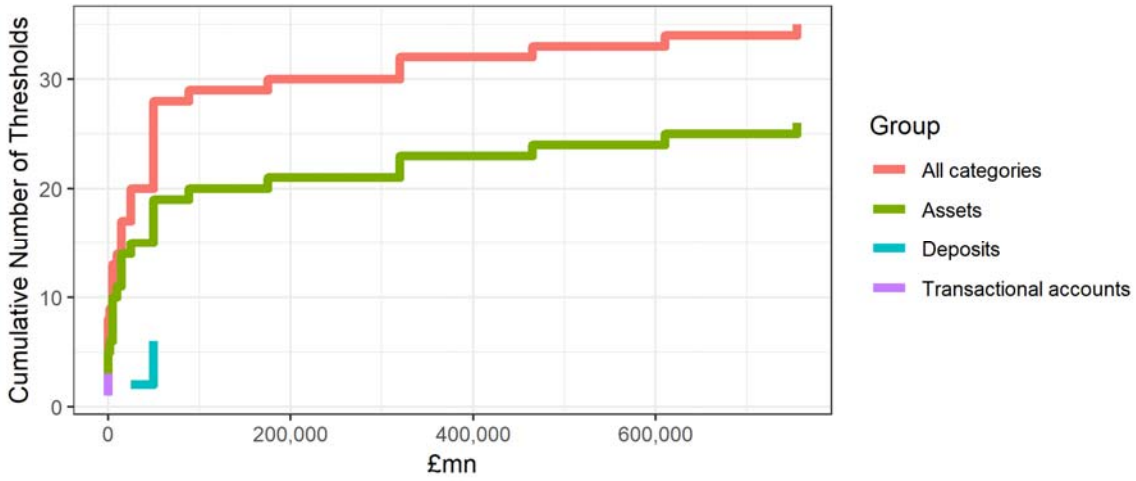

Chart 1b: Thresholds in UK prudential regulation, translated onto an assets/deposits basis up to £50b (deposit-takers)

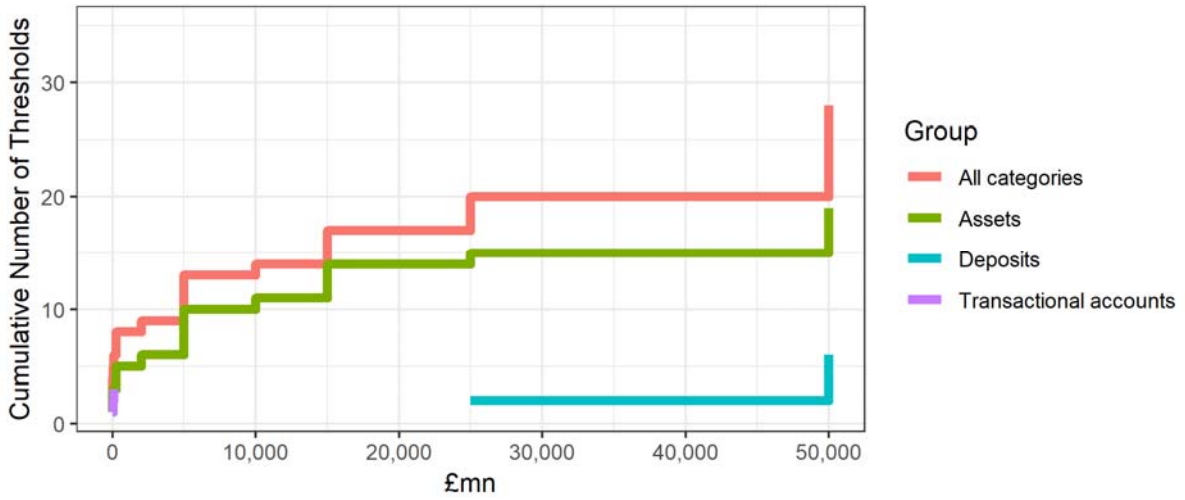

Sources: CRR, BCBS, EBA Technical Standards, PRA Rulebook, Supervisory Statements, and Statements of Policy. Note: Only includes thresholds that are either set using assets/deposits values, or can be translated into assets/deposits values. Assumes average size of transactional accounts of £500. The y-axis corresponds to the number of thresholds below the £ value for assets/deposits on the x-axis.

I'd like to know what people think about this gradient. To me, the basic gradient looks OK. It feels about right that most of the carve-outs are for firms below £50b in assets – once you reach that size, you are a fair-sized bank in a UK context and should expect to meet the great bulk of the rules. And within the £0-50b area the slope does not look particularly troublesome. But in the spirit of giving others sticks to beat us with, I invite views – particularly from small banks, who in my experience are not shy in offering us their perspectives.

But one obvious thing we can do is to map the actual distribution of banks in the UK today against the distribution of thresholds, to see how the two compare. Any sudden flattening of the cumulative distribution of banks at specific thresholds could suggest cliff-edge effects (perhaps from clusters of thresholds) that might have the potential to discourage banks from growing.

Needless to say, our tireless PRA robot has done this too. Here is what it looks like:

Chart 2a: UK prudential regulation thresholds mapped against number of banks by size (deposit-takers)

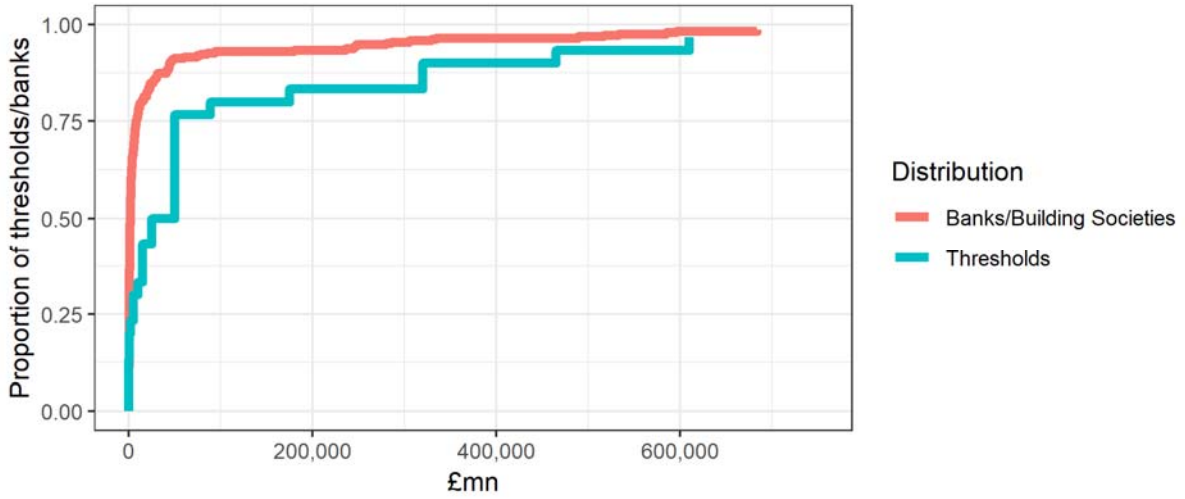

Chart 2b: UK prudential regulation thresholds mapped against number of banks by size, up to £50b (deposit-takers)

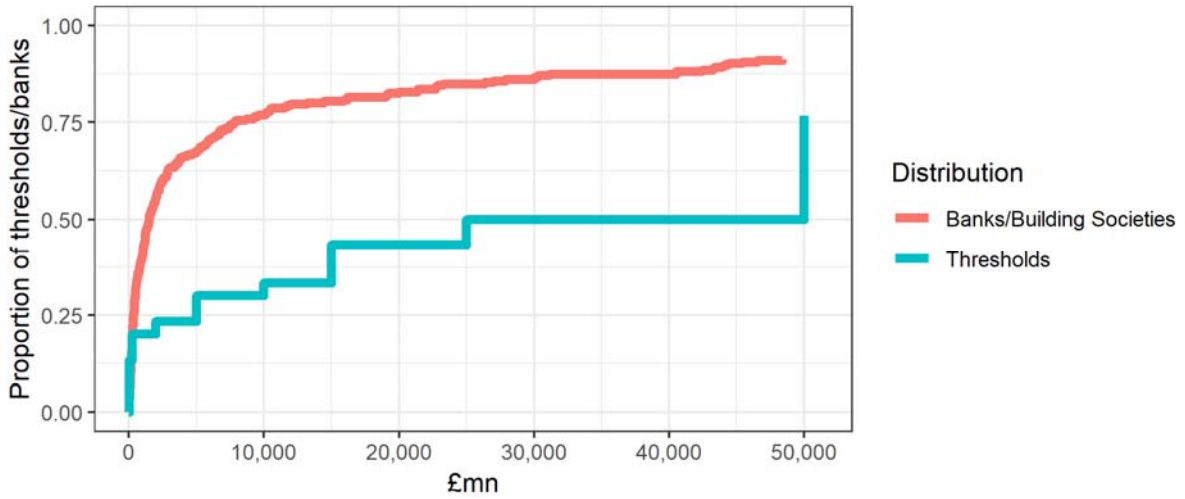

Sources: SNL Capital IQ, CRR, BCBS, PRA Rulebook, EBA Technical Standards, Supervisory Statements, and Statements of Policy. Note: Only includes thresholds that are either set using assets/deposits values, or can be translated into assets/deposits values. The y-axis corresponds to the proportion of thresholds/banks below the £ value for assets/deposits on the x-axis.

I invite views on this picture, but offer two preliminary observations myself. First, as everyone in banking knows, by number we have a concentration of firms at relatively low asset sizes. Whether this is a result of proportionality – our efforts to lower barriers to entry – or barriers to growth is a key question. Second, more novel if perhaps less fundamental: there appear to be minor discontinuities around the thresholds of £5b, £10b and £25b, and a more substantial flattening and bump at £50b. These may be worth investigating further.

Complexity

Another thing we've been looking at afresh is the complexity of our prudential rules. The scale of the global financial crisis resulted in part from the huge complexity of the financial system, and the many gaps in the rules that sought to regulate it. Fixing these gaps required more regulatory constraints and more regulatory discretion – a necessary increase in 'essential' regulatory complexity.[5] But when should we start worrying about 'too much' complexity? This is important for at least two reasons: first, excessive complexity may well be counter-productive in terms of our safety and soundness objective[6]; and second, complexity may be anti-competitive if it is harder for small firms to bear the cost of mastering it. Designing a framework to trade off costs and benefits of complexity is an emerging research question.

But there's a complementary, empirical step: we have wanted to measure the framework's complexity – to say how complex it *is*. Legal and linguistic scholarship suggests the information burden of understanding individual rules, as well as the nature of the interactions between them, both matter here.

The measurement debate tends to start with the simple metric of length: how long are our rules? Our robot has run the numbers on prudential regulation for banks and tells me that there were about 700,000 words in relevant EU rules, the PRA Rulebook (UK rules), and UK guidance (PRA supervisory statements) in 2017.

Now I've heard Andrew observe previously that that is longer than War and Peace, and others have said that this is longer than the Old Testament.[7] The inference, I think, is that our Rulebook might be a rather less good read than Tolstoy or the Bible. But actually it would not surprise me if it turned out that Andrew has our Rulebook in his living room at home and regularly dips into it for a little light refreshment following a hard day's work at the FCA.

But with our new dataset, we can explore other ways the rules have evolved from 2007 and 2017. Here are two examples.

First, as well as using a higher volume of words, we've actually used a lot more unique words. In 2007 the rules contained about 11,800 unique words. By 2017 this had increased by close to 14 percent, to over 13,400.[8] And the rules have become much more conditional. The absolute number of conditional statements

[5] *Rethinking Financial Stability* (Haldane et al, 2018)
[6] This is not a new topic, nor unique to the UK. See, for instance, Haldane paper 'The dog and the frisbee': https://www.bankofengland.co.uk/-/media/boe/files/paper/2012/the-dog-and-the-frisbee.pdf?la=en&hash=4DEAA2E6D1698A1A0891153A6B4CE70F308351D7 and the European Systemic Risk Board's paper on regulatory complexity from June 2019: https://www.esrb.europa.eu/pub/pdf/asc/esrb.asc190604_8_regulatorycomplexityquestrobustregulation~e63a7136c7.en.pdf
[7] See the van Steenis Future of Finance report: https://www.bankofengland.co.uk/-/media/boe/files/report/2019/future-of-finance-report.pdf?la=en&hash=59CEFAEF01C71AA551E7182262E933A699E952FC
[8] To put this in perspective, a typical working vocabulary for an adult native English-speaker is between 20,000 and 35,000 words. See the Economist's reference to the TestyourVocab.com test: https://www.economist.com/johnson/2013/05/29/lexical-facts

(containing words like “if” or “but”, which would tend to indicate exceptions or operational modifications) has increased, as has the relative conditionality (i.e. allowing for the fact that overall volume has gone up too).

Second, mapping networks of exposures is now a familiar tool for revealing interconnections and critical points in the financial system. Can we do the same for the system of post-crisis banking regulation?

Chart 3 shows a simplified but real representation of how four layers of regulation for UK banks interlock with each other: EU regulation in the CRR in yellow and EBA Technical Standards in red, our Rulebook in blue and Supervisory Statements in green.

This chart shows how the network of interconnections changed between 2007 and 2017. Each node is a single legal provision that explicitly cross-refers to at least one other provision. The simple observation is: the core of the network, mainly comprising the CRR, is a lot denser in 2017 than it was in 2007.

Chart 3: Network of cross-references in the regulatory framework (provisions for UK deposit-takers, 2007 and 2017)

Note: Each node represents a provision, and each edge (line between nodes) represents cross-references between provisions. For 2017, we also show the legal source for each provision: CRR (yellow nodes), Rulebook (blue), EBA Technical Standards (red) and Supervisory Statements (green). The source for the 2007 network is the FSA Handbook (only prudential provisions for deposit-takers). For visual clarity, only nodes with at least one edge are displayed, and edges on the plot are unweighted, i.e. they do not show the number of cross-references.

One way to put numbers on this is Google’s PageRank algorithm, originally designed to measure the centrality and importance of individual web pages. Chart 4 compares PageRank scores for the provisions in the 2007 and 2017 banking frameworks: you can see that connectedness increases most for the most central individual provisions. The full consequences of using and changing a handful of core rules are inherently broader in this new regulatory architecture.

Chart 4: PageRank network centrality measure (provisions for UK deposit takers, 2007 and 2017)

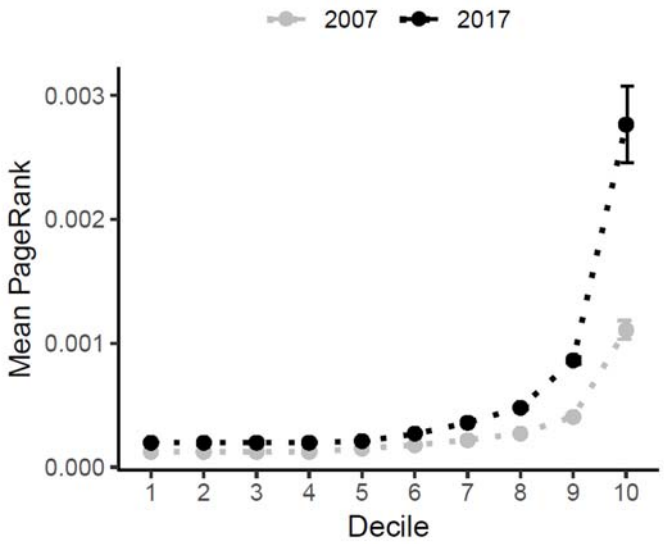

Note: PageRank summarises the centrality of a node within a network. In this case, it counts the number and quality of cross-references to a provision to estimate how important the provision is. More important, i.e. more central, provisions are likely to get more direct and indirect links from other provisions. To construct the decile plot, we calculate PageRank for each provision, rank the provisions (separately for 2007 and 2017), and then split them into ten bins (deciles). We display the mean for each bin, and the 95% confidence intervals (vertical lines).

Now, you might say: so what? Different measures give us reads on different aspects of complexity – but of course they don't tell us, in and of themselves, if it's 'good' or 'bad'. Banks and we are engaged in complicated lines of business, so unsurprisingly the regulations are complicated. There is a good deal of sense in that perspective, in my view.

Perhaps more interesting is what these new analytical capabilities can help us do as a regulator, particularly for the effectiveness of our communications. For example, they can inform policy design by telling us *ex ante* how a future change in one PRA rule could affect the interpretation of others. And they could identify rules where language could be tightened and which are frequently accessed by web users.

And we do want to look harder at the competition aspect. Small firms tell us that complexity is a real challenge and my instinct is that a simpler – simpler, not weaker – regime for small firms could benefit both our safety and soundness and competition objectives. A weaker regime for small banks would be a bad idea, particularly given the steps we have taken under our competition objective to encourage new entrants, who are by definition new and not yet tested in a downturn. But a simpler regime, perhaps borrowing from the philosophy of recent moves in Switzerland and the US, could make small firms stronger and more competitive at the same time. This is an area where we may have more room to manoeuvre following Brexit.

Action

I'm worried that this may all sound rather theoretical to some of you, and some distance from the front line of providing banking services. So let's come right back down to ground level and look at credit unions.

Credit unions provide a vital service for their members, often acting as a backstop by allowing those members to access services they might find more difficult to get from banks and building societies. The sector plays an important role in many communities – particularly in Northern Ireland.

Now credit unions, including those in Northern Ireland, may be relieved to hear that the hard-working PRA competition robot is not about to be unleashed upon them. As an aside, they may not know that they have been living with another of our robots for some time. A few years back some of our economists from the wider Bank of England teamed up with our credit union supervisors and found out that by using some different analytical techniques to study regulatory returns we could do a better job of predicting which credit unions would get into trouble next.[9] This is highly relevant given we oversee around 440 credit unions with a team of 9 staff. They created a "Risk Profiler" bot which is not perfect but plays a useful role in directing our enquiries.

More recently, our credit union supervisors got in touch with our competition team about working on barriers to growth issues. They observed that the structure of our requirements creates some difficult cliff-edge effects for growing firms.

This arises in the following way. For all deposit-takers in the UK (banks, building societies and credit unions) capital requirements go up in tiers as they get larger. This works as a "slab" system. To illustrate this let me take the perilous step of drawing an analogy with tax, perilous in that it is considered a cardinal sin for central bank officials to utter the word "tax". You can contrast this "slab" system with income tax. When a firm's size exceeds a certain threshold, it has to carry an extra layer of capital not only for each asset it takes on above that threshold, but for its entire asset base. To make the comparison with income tax, it's as though if your income exceeded £50,000, the higher tax rate of 40% was then applied not only to your earnings above £50,000 but also to all of your earnings below the threshold.

We tend to think – in line with the international consensus in Basel – that this tougher "slab" approach is a good thing for deposit-takers because it should discourage them from becoming too large too quickly. I think that makes sense for banks, but we've come to the view that it is really not relevant for credit unions in the UK. We also want to avoid discouraging growing firms, and are confident of our ability to manage the failure of smaller credit unions smoothly with our colleagues at the FSCS. Therefore we are launching a

[9] For more background on this work, see https://www.bankofengland.co.uk/working-paper/2017/the-determinants-of-uk-credit-union-failure

consultation today in which we propose to change capital requirements to simplify them, remove barriers to growth and deliver a greater degree of financial resilience for smaller credit unions that are more likely to fail.

We propose to:

- get rid of our “slab” system and replace it with a graduated, income-tax-style approach where larger credit unions would be required to hold 5% capital for assets up to £10m, 8% for assets between £10m and £50m and 10% for assets above £50m, as opposed to today’s system where firms have suddenly to hold 10% on their entire asset base if they grow a single pound over £10m in assets;

- maintain capital requirements at 3% for credit unions with total assets less than £5m, but introduce a “monitoring zone” of 3-5% whereby if a smaller credit union’s capital falls below 5% we would step up our supervisory activity. In time we hope this will help contain the rate credit union failures, which is currently disproportionately high amongst these smaller firms; and

- simplify the regime by eliminating other capital requirements thresholds linked to non-asset metrics such as membership numbers and variety of business lines.

The effect of these changes should be to reduce overall capital requirements for the credit union sector by roughly one quarter, and encourage the growth of successful credit unions by removing several cliff-edges present in our current regime.

Conclusion

We at the Bank have been clear that we have no appetite to weaken the level of resilience in the financial system. But this does not mean that we should stand still, particularly when we have opportunities to advance safety and soundness and competition at the same time. The changes we are proposing today are an example of what we can do when we bring a bit of fresh thinking, supported by our emerging army of regulatory robots, to the question of barriers to growth.