

Andrew G Haldane: On microscopes and telescopes

Speech by Mr Andrew G Haldane, Executive Director and Chief Economist of the Bank of England, at the Lorentz centre workshop on socio-economic complexity, Leiden, Netherlands, 27 March 2015.

* * *

The views are not necessarily those of the Bank of England or the Monetary Policy Committee. I would like to thank Thibaud de Barman, Lucy Canham, Chiranjit Chakraborty, Jeremy Franklin, Rashmi Harimohan, John Hill, Bradley Hudd, Ben Kempley, Lu Liu, Katie Low, Anna Jernova, Carsten Jung, Damien Lynch, Yaacov Mutnikas, Tobi Neumann, Paul Robinson and David Ronicle for their assistance in preparing the text. I would also like to thank Stephen Burgess, Oliver Burrows, Pavel Chichkanov, Nicholas Fawcett, Bob Hills, Ronan Hodge, Glenn Hoggarth, Catrin Jones, Vas Madouros and Simon Whitaker for their comments and contributions.

Accompanying charts and the table can be found on the [Bank of England's website](#).

At least since the financial crisis, there has been increasing interest in using complexity theory to make sense of the dynamics of economic and financial systems (Newman (2011), Arthur (2014)). Particular attention has focussed on the use of network theory to understand the non-linear behaviour of the financial system in situations of stress (Gai and Kapadia (2011), Haldane and May (2011), Gai, Haldane and Kapadia (2011)). The language of complexity theory – tipping points, feedback, discontinuities, fat tails – has entered the financial and regulatory lexicon.

Some progress has also been made in using these models to help design and calibrate post-crisis regulatory policy. As one example, epidemiological models have been used to understand and calibrate regulatory capital standards for the largest, most interconnected banks – the so-called “super-spreaders” (Craig et al (2014)). They have also been used to understand the impact of central clearing of derivatives contracts, instabilities in payments systems and policies which set minimum collateral haircuts on securities financing transactions (Haldane (2009)).

Rather less attention so far, however, has been placed on using complexity theory to understand the overall *architecture* of public policy – how the various pieces of the policy jigsaw fit together as a whole. This is a potentially promising avenue. The financial crisis has led to a fundamental reshaping of the macro-financial policy architecture. In some areas, regulatory foundations have been fortified – for example, in the *micro-prudential* regulation of individual financial firms. In others, a whole new layer of policy has been added – for example, in *macro-prudential* regulation to safeguard the financial system as a whole (Hanson, Kashyap and Stein (2010)).

This new policy architecture is largely untried, untested and unmodelled. This has spawned a whole raft of new, largely untouched, public policy questions. Why do we need both the micro- and macro-prudential policy layers? How do these regulatory layers interact with each other and with monetary policy? And how do these policies interact at a global level? Answering these questions is a research agenda in its own right. Without answering those questions, I wish to argue that complexity theory might be a useful lens through which to begin exploring them. The architecture of complex systems may be a powerful analytical device for understanding and shaping the new architecture of macro-financial policy.

Modern economic and financial systems are not classic complex, adaptive networks. Rather, they are perhaps better characterised as a complex, adaptive “system of systems” (Gorod et al (2014)). In other words, global economic and financial systems comprise a nested set of sub-systems, each one themselves a complex web. Understanding these complex sub-systems, and their interaction, is crucial for effective systemic risk monitoring and management.

This “system of systems” perspective is a new way of understanding the multi-layered policy architecture which has emerged since the crisis. Regulating a complex system of systems

calls for a multiple set of tools operating at different levels of resolution: on individual entities – the *microscopic* or *micro-prudential* layer; on national financial systems and economies – the *macroscopic* or *macro-prudential* and *monetary policy* layer; and on the global financial and economic system – the *telescopic* or *global financial architecture* layer.

The architecture of a complex system of systems means that policies with varying degrees of magnification are necessary to understand and moderate fluctuations. It also means that taking account of interactions between these layers is important when gauging risk. For example, the crisis laid bare the costs of ignoring systemic risk when setting micro-prudential policy. It also highlighted the costs of ignoring the role of macro-prudential policy in managing these risks. That is why the post-crisis policy architecture has sought to fill these gaps. New institutional means have also been found to improve the integration of these micro-prudential, macro-prudential, macro-economic and global perspectives. In the UK, the first three are now housed under one roof at the Bank of England.

In what follows, I first set out some background on the dynamics of a complex system of systems using some stylised examples. I then discuss some stylised facts on the “system of systems” that is today’s economic and financial network. Finally, I draw out some tentative conclusions for future research and policy which follow from viewing the macro-financial system through this lens.

The architecture of complexity

The literature on complexity theory, and its implication for system dynamics, is now deep and rich (Newman (2011)). Although there is no generally-accepted definition of complexity, the one contained in Herbert Simon’s classic 1962 article on the *Architecture of Complexity* – “one made up of a large number of parts that interact in a non-simple way” – continues to capture well its everyday essence (Simon (1962)). In complex systems, the whole behaves very differently than the sum of its parts.

Although there is no single unifying theory of complexity, the dynamic properties of complex systems are now reasonably well-understood, based on analytical and experimental studies of networks of all types – physical, natural, social, biological and economic (Ladyman et al (2013)). These dynamic properties include non-linearity; discontinuities in responses to shocks; amplifying feedback effects; and so-called “emergent” system-wide behaviour which is difficult to predict from the behaviour of any one element.

These properties of complex systems typically give rise to irregular, and often highly non-normal, statistical distributions for these systems over time. This manifests itself as much fatter tails than a normal distribution would suggest. In other words, system-wide interactions and feedbacks generate a much higher probability of catastrophic events than Gaussian distributions would imply (Newman (2005), Gabaix (2009)). They may also result in distributions which are multi-modal, consistent with models of multiple equilibria (Bisin et al (2011)).

The topology and wiring of these complex systems appears, perhaps predictably, to have a crucial bearing on their resilience to shocks. Many complex networks have been found, in practice, to exhibit a “scale-free” property (Barabasi and Albert (1999)). That is to say, they comprise a core set of nodes with a large number of connections and a large set of peripheral nodes with few connections. There is a core-periphery, or hub-and-spokes, network configuration. This scale-free property has been found in everything from food webs to the World Wide Web, from eco-systems to economic systems, from synapses to cities, from social networks to financial networks (Jackson (2010)).

These scale-free topologies have important, if subtle, implications for system resilience. For example, core-periphery models have been found to be very robust, at a systemic level, to random shocks. That is because these shocks are very likely to fall on peripheral nodes unconnected with, and hence unlikely to cascade through, the system as a whole. But these

systems are also vulnerable to targeted attack on the core nodes – the “super-spreaders” – whose hyper-connectivity risks generating a systemic cascade (Albert, Jeong and Barabasi (2000)).

Another typical feature of complex systems is that they tend to organise themselves as a hierarchy, with a well-defined structure of systems and sub-systems (Simon (1962, 1976)). Herbert Simon believed hierarchical structures of a particular type were likely to dominate, namely ones which were “decomposable”. By this he meant organisational structures which could be partitioned such that the resilience of the system as a whole was not reliant on any one sub-element.¹ For evolutionary reasons of survival of the fittest, Simon posited that decomposable networks were more resilient and hence more likely to proliferate as a species (Simon (1962)).²

While Simon’s evolutionary theory may be a reasonable long-run description of some real-world complex systems – natural and biological – it may be less good as a description of the evolution of *socio-economic* systems. The efficiency of many of these networks relies on their hyper-connectivity. There are, in the language of economics, significantly increasing returns to scale and scope in a network industry. These returns increase with network connectivity (Goldin and Mariathason (2014)). Think of the benefits of global supply chains and global interbank networks for trade and financial risk-sharing. This provides a powerful secular incentive for non-decomposable socio-economic systems.

Moreover, if these hyper-connected networks do face systemic threat, they are often able to adapt in ways which avoids extinction. For example, the risk of social, economic or financial disorder will typically lead to an adaptation of policies to prevent systemic collapse. These adaptive policy responses may preserve otherwise-fragile socio-economic topologies. They may even further encourage the growth of connectivity and complexity of these networks. For example, policies to support “super-spreader” banks in a crisis may encourage them to become larger, and more complex, still over time (Haldane (2009)). The combination of network economies, and policy responses to failure, means socio-economic systems may be less Darwinian, and hence decomposable, than natural and biological systems.

It is against this backdrop that a complex, socio-economic “system of systems” may emerge. This can be defined as one comprising an interlocking set of individually complex webs ((Gorod et al (2014)). The system of system concept initially emerged for engineering and enterprise systems, which involved the multi-layered assembly of component parts. But it has since found its way into a number of other domains including military planning, ecological evolution, power grids, transport networks and neurological structures (Gao et al (2014)).

Although still in its infancy, there are some general properties of a “system of systems” perspective that are worth bringing out. For example, Kurant and Thiran (2006) look at the behaviour of a particular topology – a *layered* complex network. Specifically, they focus on the behaviour of transport networks with a two-layer structure. Simulations of this network suggest that monitoring risk on a layer-by-layer basis is likely to understate significantly the risk facing each individual layer.

Layered complex networks may also be less robust to failure than might be apparent from assessing the resilience of each layer in turn. In other words, the risks in a layered network are strikingly different than the sum of their parts. The greater the complexity of each layer, and the stronger the correlation between layers, the greater is this vulnerability (Kurant and

¹ As an example, early TV sets were built in a non-decomposable way, which made them vulnerable to the failure of one element. Later, TV sets, cars and other complex gadgets tend to be constructed in a “decomposable” fashion to improve their resilience.

² Charles Perrow’s concept of “tightly coupled” – non-decomposable – systems is closely linked (Perrow (1984)).

Thiran (2006)). In some respects, this is the counterpart of Simon’s “decomposability” hypothesis in a system of systems context.

Most recently, research has focussed on the controllability of complex, layered networks (Liu, Slotine and Barabasi (2011), Gao, Liu, D’Souza and Barabasi (2014)). It has tended to find that, even when the dimensionality of a network is large (a large number of layers), effective control can be exercised by acting on a relatively small number of key layers or nodes. This is particularly the case when the network has scale-free properties – in other words, a core-periphery-type topology.

These points can be brought to life using some simulations of the statistical distribution of a simple set of systems of systems. These are constructed by mixing together component distributions, which proxy the layers or sub-systems. These layers may themselves be complex. This mixing of layers, or distributions, is done using an assumed distribution of correlations, which may itself be non-normal and complex. In other words, there are multiple layers of complexity in this system of systems. The inputs to the simulations are shown in Table 1.

Chart 1 shows a set of joint (system-wide) distributions from these simulations, where each line represents a probability contour of the distribution.³ Chart 1a is the baseline case. It involves a mix of two normal distributions with a low correlation coefficient (0.3) where these correlations are themselves normally distributed.⁴ This joint distribution represents a two-layer system of systems, with neither layer complex (hence the normal distribution for each) and where the correlation between the sub-systems is weak and regular. The resulting joint distribution is slightly more elliptical than the normal – meaning a greater likelihood of large positive or negative outcomes occurring simultaneously – but is otherwise unexceptional.⁵

We can now add progressively greater degrees of complexity to this base case to assess its impact on the distribution of systemic risk. Chart 1b raises the correlation between the two layers to 0.8; there is now strong feedback between the sub-systems. It results in a notable elongation of the system-wide distribution and an even greater probability of good or bad news striking simultaneously; it becomes more non-normal, closer to the statistical properties we expect from natural, biological and social networks.

If we allow one of the layers of this cake to exhibit the properties of a complex system, by assuming it is t- rather than normally-distributed, then the system-wide distribution is shown in Chart 1c.⁶ The tails of this distribution have now further widened. And if we allow for a non-normal distribution of correlations between the sub-systems – by assuming correlations rise when there is a bad draw, as during crises – the lower tail is larger still and the mass is skewed downwards (Chart 1d).

The final, and perhaps most interesting, case is shown in Chart 1e. This shows a three-layer network. Two of these layers are themselves complex (they are t-distributed) and all of the layers are strongly correlated in the network tree. The resulting system-wide distribution is highly irregular. It is also heavily fat-tailed, meaning that catastrophically good or bad outcomes are now much more probable than the normal distribution would suggest. Its topology is significantly more complex than any of its individual layers.

These simulations, although simple, provide some insight into the likely behaviour of a complex system of systems. For example, they suggest that viewing risk through the lens of a single layer is likely to provide a significantly distorted picture of the true risk distribution,

³ Which are based on one million replications of the data generating process.

⁴ Technically, this is done using the so-called Gaussian copula.

⁵ With zero correlation between the two distributions, Chart 1a would be a set of concentric circles.

⁶ The t-distribution has fatter tails than the normal. In the simulations, we assume 5 degrees of freedom.

with the probability of tail events materially under-estimated. As an example of that, Chart 1f looks at the unconditional and conditional distributions of one of the sub-systems in Chart 1c. The conditional distribution is conditioned on outcomes in the second layer lying in the lower half of the distribution.

A risk manager looking at the unconditional distribution (a single layer perspective) would significantly under-estimate the true tail risk they were facing. This has direct parallels with the pre-crisis situation, when the actions of individual risk managers at banks, and individual supervisors of those banks, ignored the role of systemic risk when gauging individual firm risk. This led to a material under-estimation of individual firm risk. The micro-prudential microscope was wrongly focussed.

Ignoring a layer, or sub-system, is likely to be equally distorting. That can be seen by comparing the risk distributions in Chart 1d (two layers) and 1e (three layers). If the world involves three layers of complexity, then missing a layer will lead to a significant under-estimation of risk. This, too, has a parallel with the pre-crisis situation where the macro-prudential layer was essentially ignored by both micro-prudential and macro-economic policymakers. This led both to under-estimate the policy risk they faced.

Pulling this together, what are the public policy implications which follow from this complex system of systems perspective? First, it underscores the importance of accurate data, and timely mapping, of each layers in a system of systems. This is especially important when these layers are themselves complex. Granular data is needed to capture the interactions within and between these complex sub-systems.

Second, modelling of each of these layers, and their interaction with other layers, is likely to be important, both for understanding system risks and dynamics and for calibrating potential policy responses to them.

Third, in controlling these risks, something akin to the Tinbergen Rule is likely to apply. There is likely to need to be at least as many policy instruments as there are complex sub-components of a system of systems if risk is to be monitored and managed effectively. Put differently, an under-identified complex system of systems is likely to result in a loss of control, both system-wide and for each of the layers.

The architecture of macro-financial systems

How, then, does this theory relate to real-world, macro-financial systems? These systems are likely to contain many moving parts. Moreover, these moving parts are likely to be significantly more tightly-coupled than in the past as a result of financial and global integration. In other words, the global economic and financial system may, over recent decades, have become a “system of systems”, with multiple, interacting layers each a complex system in its own right.

Chart 2 provides a stylised characterisation of those layers, decomposed four ways. At the highest resolution – the “micro-prudential” layer – are individual financial firms. These are, if you like, the atoms of the financial system. Like atoms, however, some individual banks are themselves complex entities, with many moving and interacting business parts.

At one lower level of resolution – the “macro-prudential” layer – is the financial system. This comprises interactions between financial firms in the network, as might arise from counterparty relationships in interbank, repo and derivatives markets. This layer is akin to an organ, like the brain, whose behaviour is the result of interactions between complex neurological sub-components.

At a lower level of resolution still – the “macro-economic” layer – is the national economy. This comprises complex interactions between the financial sector and the wider economy – the flows of funds which intermediates money from owners to borrowers. It is akin to the

physiology of a human, whose behaviour in the result of complex interactions among its organs, themselves complex sub-components.

Finally, at the lowest level of resolution of all – the “telescopic” layer – is the global economic and financial system. This involves cross-border trade and financial interactions between countries, with flows of goods, services or information at increasing volumes and velocities. It is akin to interactions within a social network or across the World Wide Web.

At least in principle, each of these individual layers could be complex. These layers are also likely to interact. For example, there are likely to be strong interactions between the financial system and the macro-economy and between individual national financial systems and the global financial system. This coupling between sub-systems adds to the degree of complexity of the system as a whole.

Let me present some stylised facts on each of these layers, from microscopic through telescopic, to illustrate why each has the characteristics of a complex sub-component. I will also present some evidence on interactions between these sub-systems, which means we have a genuine macro-financial system of systems.

(a) Complexity among individual firms

There is no off-the-shelf measure of organisational complexity, whether within banks or more generally, but a number of proxies have recently been constructed. For example, the Basel Committee has recently devised a range of complexity metrics for the world’s so-called SIFIs – systemically important financial institutions.⁷ So far, 30 institutions globally have been designated by the Financial Stability Board (FSB) as SIFIs, from across a range of countries.

Chart 3 looks at a set of four complexity proxies for each of the designated SIFIs in 2006, the year before the financial crisis broke: total balance sheet size; the notional value of their derivatives portfolio; the number of legal entities in the group; and their trading assets. They are shown by their country of incorporation.

In 2006, the average SIFI had a balance sheet of \$1.35 trillion – roughly the same as the annual GDP of a small G7 country. The largest SIFI had a balance sheet totalling \$3.0 trillion. This is many multiples of the largest non-financial firm. It also only covers on-balance sheet positions. Turning to off balance sheet positions, the derivatives portfolio of an average SIFI totalled \$19 trillion in notional terms in 2006.⁸ For the largest SIFI, it was \$60 trillion. This, too, is many multiples of the largest non-financial firm’s position.

These massive SIFI balance sheets were, in turn, spread across a very large array of distinct legal entities within each group. In 2006, the average SIFI had around 328 legal entities within the group. The largest had almost 900 distinct legal entities. Most of the largest global SIFIs had numbers of legal entities running to four figures. By comparison, the world’s largest non-financial firms had less than half this number of legal affiliates (Carmassi and Herring (2014)).

A final complexity metric is the proportion of so-called trading assets. This is a diagnostic on the diversity of a bank’s business model – whether it is engaged in investment as well as commercial banking. Some of these trading assets will also be complex and hence difficult to price and trade – so-called Level 2 and 3 assets. As Chart 3 shows, the average SIFI had a trading book of around a quarter of assets in 2006. For the largest, this was closer to around 50% of assets.

⁷ See BCBS (2014).

⁸ Notional exposures are likely to be a poor proxy of the underlying risk on this portfolio, but provide an indication of the scale of gross derivative positions.

Since the crisis, there is little evidence of these complexities having reversed. Chart 4 updates the four metrics using end-2013 data. The balance sheet of the average SIFI has not shrunk. In fact, it had risen to around \$1.8 trillion by end-2013. For the largest SIFI, it was around \$4.0 trillion. The pattern is even more striking off balance sheet. For the average SIFI, the derivatives book has risen by over 50% to over \$30 trillion by 2013. For the largest, it had risen to \$75 trillion.

On the other complexity metrics, the number of legal entities for the average SIFIs is essentially unchanged comparing 2006 and 2013. And despite a significant rise in risk weights, there had been only a modest slimming of SIFI trading books by 2013, with the average share of assets going from 22% to 19%.

An alternative, more micro-economic, approach to gauging organisational complexity is to look at internal structures. For example, Collinson and Jay (2012) use a five-way classification of complexity: governance; internal structures; staff capabilities; roles and responsibilities; and corporate culture. They use this to evaluate the world's largest 200 companies, including 26 banks.

Financial services rank sixth overall as an industry on complexity grounds. Within this, however, there is a striking diversity in performance, with a number of banks standing out as having high levels of complexity and low performance. Chart 5 plots banks' complexity and performance, based on a study by Simplicity Consulting in 2012.⁹ It shows an inverted U-shaped pattern. While modest degrees of bank complexity are good for performance, beyond a threshold these complexity benefits go into reverse and performance worsens.

This U-shaped pattern is consistent with micro-economic theory. It suggests economies of scale and scope in banking can give way to diseconomies beyond some threshold. There is empirical evidence suggesting diseconomies of scale may operate for some SIFIs, consistent with them being "too complex to manage" (Davies and Tracey (2014)). Taken together, this evidence suggests the world's SIFIs are, and remain, highly complex entities in their own right.

(b) Complexity within the financial system

As with individual banks, there are no perfect metrics of financial system complexity. But maps of the topology or wiring of the financial sector can be revealing about the strength, direction and complexity of the counterparty links between, and interactions among, financial firms. As one example, Chart 6 looks at the interbank connections between UK banks in 2013.¹⁰

This suggests the wiring of the interbank network is messy and dense. The network's epicentre is a small set of large banks with a very large number of connections – "super-spreader" banks. There are, in addition, a large number of peripheral banks with few connections. In other words, the interbank system has all the properties of a classic scale-free network.

Chart 7 looks at the pattern of interbank funding exposures among UK banks. It, too, suggests a complex web with scale-free properties. Networks with these properties are prone to failures among the super-spreader nodes. Since these are likely to be the entities most exposed to "too complex to manage" problems, this topology harbours a potential systemic vulnerability, the type of which was exposed in 2008.

⁹ "True Cost of complexity in the banking sector", Simplicity Consulting (2012).

¹⁰ It is taken from Langfield, Liu and Ota (2014). Exposures comprise unsecured loans, marketable securities, CDS sold minus CDS bought, securities lending and repo and derivatives exposures.

Complex interactions are not confined to the banking sector. The non-bank (“shadow banking”) sector has become increasingly important over recent years. In the US, estimates suggest it was broadly equivalent in scale to banking in the run-up to the crisis. At a global level, non-bank intermediation is currently estimated at around \$75 trillion, around half of the assets of the global banking system.

Because many of these entities sit outside of the regulatory perimeter, data on their exposures and interactions is thin on the ground. But work in the US to map these interactions suggests they too are dense and complex (Pozsar et al (2010)). Chart 8 presents a stylised map of the shadow banking sector and its interaction with the banking sector. These interactions are highly complex and multi-layered.

(c) Complexity of macro-financial interactions

The links and interactions between the financial sector and the wider economy are, in some respects, even less well mapped than the shadow banking sector. They, too, have evolved significantly over recent decades in scale and complexity.

As an example of that evolution, Chart 9 looks at a representation of the flow of funds between various sectors in the UK – companies, households, the rest of the world and government. It also shows the role of the financial sector in intermediating these flows of funds between ultimate asset-owners (shown at the top) and liability-holders (shown at the bottom). The intermediation sector is further decomposed into banking and “other” financial intermediaries such as pension, hedge and other funds and insurance companies. The Bank of England’s is also shown.

These balance sheets are scaled relative to GDP at two dates, 1978 and 2011. The changes over time in the scale and pattern of intermediation are striking. In 1978, the UK banking system was dominated by foreign banks. Banks’ assets summed to a little over 100% of GDP. Non-bank intermediaries, meanwhile, accounted for little more than 50% of GDP and were dominated by traditional real-money investors, such as pension funds and insurance companies. The Bank of England’s balance sheet was a mere 6% of GDP.

Fast forwarding to 2011, this flow of funds picture has been transformed. The balance sheets of ultimate asset-owners/liability-holders has risen from around 200% of GDP to around 600% of GDP – a threefold rise. The scale and pattern of intermediation between these sectors has also been transformed. The banking sector has also risen substantially over this period. UK universal banks now dominate the landscape, accounting for 80% of total banking assets. They are the new “super-spreaders” of the UK financial sector.

Non-bank intermediation has also seen a dramatic scaling-up of its balance sheets, rising sixfold to 300% of GDP. Within this, growth has been particularly strong within the non-traditional sectors (such as hedge funds and unauthorised funds) and among new types of vehicle (such as ETFs and private equity funds). These are a new strain of potential “super-spreader”. The Bank of England’s balance sheet has also grown rapidly, rising to over 25% of GDP, largely courtesy of QE.

Overall, the picture is of a much larger, much more diverse and potentially considerably more complex set of channels of intermediation than in the past. Of course, this stylised map offers only the most cursory of indications of the true degree of complexity in individual interactions within and between sectors.

(d) Complexity of global macro-financial interactions

A final dimension to network complexity is the global one. There has been a dramatic rise in the scale of global trade and, in particular, capital market integration over the past few decades. Both are now at their highest levels for several centuries and, most probably, ever (Chart 10).

For global trade, Chart 11 compares trade flows across 70 countries constituting around 80% of world exports at two dates – 1995 and 2013. The size of the nodes reflects the importance of a country to world trade and the thickness of the links the size of trade flows. It is clear that the wiring of the international trading systems has become considerably more dense and complex over this period, reflecting the deepening and lengthening of global value chains (OECD (2013), de Backer and Miroudet (2014)).

This highly interconnected and complex set of global supply chains has the potential to propagate disturbances much more virulently than in the past, as two recent events illustrate. The first was the correlated collapse of world trade during the financial crisis. Between 2008Q1 and 2009Q1, real world trade fell by almost 15%. Correlations between regional import volumes rose to close to unity at the height of the crisis (Chart 12). The strength and length of global supply chains caused a synchronous collapse in trade (Bems et al (2012)).

The second was the Japanese earthquake and tsunami in 2011. Because Japan plays a central role in global supply chains as a producer of high-value intermediate goods, this had a cascading effect on global production. For example, as Japanese car producers ceased production, their foreign affiliates were unable to compensate because they depending on inputs from Japan (OECD (2013)). Deepening global supply chains mean a large number of countries are now potential “super-spreaders” of global trade contagion.

The story is similar, if on an even more dramatic scale, when it comes to global banking. Chart 13 plots the network of cross-border bank lending flows across a wide range of countries, in 1980 and in 2007. As with trade, the wiring of the international banking network has become significantly more dense and complex (Minoiu and Reyes (2013)).¹¹ And as with domestic banking networks, there is a clear core-periphery pattern to the network, with a small number of nodal “super-spreader” banking sectors serving as international lenders to the periphery banks.

The global financial crisis provided a telling illustration of the susceptibility of such a network to a weakening, or failure, of the core banking sector. As these core sectors came under stress, it caused a significant cascading of global interbank contagion. Cross-border interbank lending flows fell by around 20% between 2008Q2 and 2010Q2. The largest cross-border banks – the “super-spreaders” – appear to have played a central role in instigating and propagating this global contagion (Hills and Hoggarth (2013)).

(e) Correlations among sub-systems

Another dimension when gauging risks to the macro-financial system is the correlation *between* layers of the system. For correlations between individual banks and the financial system as a whole, one metric is provided by looking at the common component within the joint distribution of bank equity prices. For example, the first principal component of UK bank equity prices accounts for around 70% of the overall movement in bank equities. At the time of the financial crisis, this common component rose to nearer 90%. These correlations are high and highly crisis-sensitive, suggesting a close link between individual and systemic bank risk.

The picture is similar when looking at correlations between developments in the financial system and the wider economy. For example, the pairwise correlation between the financial and non-financial components of UK equity prices is high, at around 0.8–0.9. The picture is similar if we move from the national to the international level: the first principal component of international equity prices accounts for around 80–90% of the variation. Each of these

¹¹ This analysis uses the Bank for International Settlements bilateral locational statistics. This does not capture flows between periphery countries, so there is a risk that it overstates the importance of the core. See Minoiu and Reyes (2013) for more information

correlations is, in their own way, a rather imperfect metric of cross-system interaction. Nonetheless, they suggest these interactions are, in general, both important and strong.

A final, indirect and reduced-form, piece of evidence on the complexity of the macro-financial system comes from looking at the statistical distribution of behaviour in these systems. For other complex systems, physical, natural, biological and social systems, this statistical distribution is known to be highly non-normal (Newman (2005)). Indeed, many of these systems are found to be Power-law distributed, with much fatter tails than would be expected from a normal distribution.

To what extent are these statistical properties evident in macro-economic and financial time-series data? Chart 14 looks at the statistical distribution of GDP growth, credit growth, equity prices and rice prices, using a long time-series covering several hundreds of year. These densities have been transformed such that, if the underlying distribution is normal, the dots trace a straight line. Any deviation at the extremities is evidence of a fat-tailed distribution.

It is clear from Chart 14 that all four series exhibit a significant degree of fat-tailedness, considerably more so than a normal distribution. This is consistent with economic and financial systems, locally and globally, exhibiting complex system of systems properties. For example, a four-standard deviation event – a true catastrophe – would under a normal distribution be expected to occur roughly every 15,000 years. Under the estimated distributions for economic and financial systems, such an event would occur every 10 to 15 years.

Implications for policy design

So what might this mean for the design of the policy architecture? Let me mention three areas where clear progress has been made since the crisis, but where more may be needed in the period ahead if we are to manage effectively this complex, macro-financial system of systems: *data*, *modelling* and *policy design*.

(a) Data

Monitoring risk in a system of systems requires data on each of the sub-systems, often at a high level of granularity. Often in response to crisis, efforts have been made over time to improve data on economic and financial systems. For example, there have been successive waves of improvement to statistics on international banking flows, starting in the 1970s. In response to the latest crisis, renewed efforts have been made to fill data gaps and to improve the degree of disclosure by banks of their balance sheet positions (FSB (2012)).

As one example, data on banks' cross-border assets and liabilities only distinguish between banks and non-banks. In practice, a richer sectoral decomposition of these data is often desirable to understand interactions between different parts of the financial sector and between it and the wider economy (Hills and Hoggarth (2013), Lane (2014)). For example, the dollar funding crisis which faced European banks in 2011/12 involved a complex string of cross-border financing. Funds raised in the US from money market funds were round-tripped back into the US through the purchase of asset-backed securities by European banks. When US money funds withdrew their dollar funding, this chain went into reverse gear causing a dollar funding problem. Understanding those links in the cross-border financing chain is impossible without detailed sectoral data. That is now in train.

Yet there remain significant data constraints on our ability to map the macro-financial system of systems in the detail necessary to capture accurately interactions within and across sub-systems. Even among the world's largest banks, data on their bilateral exposures to one another remains partial and patchy, especially for off balance sheet positions and securities holdings. That means large parts of the core of the international banking map remain, essentially, uncharted territory.

These data gaps are even more acute when moving beyond the banking system. Large parts of the non-bank sector remain in the shadows from a data perspective. For example, reliable measures of the aggregate and/or individual leverage, maturity and foreign currency mismatch positions of non-bank intermediaries – some of the key vulnerability metrics – are thin on the ground. The FSB has an agenda to fill these data gaps.

More broadly, flows of funds data, both within the financial sector and between it and the wider economy, domestically and internationally, is at present partial. Chart 15 takes the UK flows of funds map shown earlier and colour codes this according to one measure of its data availability. There remain a number of key areas for improvement, particularly within the non-banks part of the intermediation chain, as well as among companies, households and externally.

In response to this, the Bank and the UK's Office for National Statistics (ONS) have recently initiated a project to improve the UK's flows of funds data. Although at an early stage, this aims to provide more granular picture of sectoral ("who-to-whom") flows. Once complete, this should enable interactions between and within the financial and real sectors of the UK economy to be better tracked and modelled.

This pattern of data gaps is mirrored globally. Global flow of funds data beyond banking is patchy. Data on flows, and in particular stocks, of cross-border portfolio and foreign direct investment are often low in quality and lack timeliness. For example, official data suggest the UK has a net external *liability* position of around 20–30% of GDP. But if foreign direct investment stocks are valued at market (rather than book) prices, this switches to a net *asset* position of around the same amount (Bank of England (2014)).

Derivatives positions and net foreign currency positions are also imperfectly captured in global flows of funds data, particularly for companies. For example, over recent years many emerging market companies have borrowed cheaply in dollar-denominated instruments. But gauging the scale of that borrowing, and whether it is hedged either with derivatives or dollar income receipts, is a real data challenge. This is a significant gap because unhedged dollar borrowing is one potential fault-line in the international financial system at present (BIS (2014)).

(b) Modelling

Once the necessary data is in place, a natural next question is how best to use this to understand the dynamics of the macro-financial system of systems. There are a number of potentially fruitful, if fledgling, avenues of research currently being pursued on that front.

In making sense of the macro-financial system, it would be desirable to have a quantitative framework for understanding and evaluating interactions both within the financial system and between it and the wider economy. Pre-crisis, such models were close to non-existent. Indeed, many mainstream macro-economic models did not even contain a well-defined financial sector, much less interactions within it (Roger and Vleck (2011)).

The Bank of England was an early pioneer of an approach which placed the financial sector centre-stage, with its so-called RAMSI (Risk Assessment Model of Systemic Institutions) model (Burrows et al (2012)). A stylised overview of RAMSI is given in Chart 16. RAMSI embodies interactions between the elements of individual banks' balance sheets (the micro-prudential layer), between individual banks (the macro-prudential layer) and between the financial system and the economy (the macro-economic layer).

As an example, RAMSI embodies an interbank network which allows counterparty network contagion to propagate. It also contains a behavioural model of funding markets and banks' liquidity positions which can generate liquidity crisis externalities (Aikman et al (2009)). There are also explicit behavioural links between macro-economic variables and bank balance sheets (through arrears, defaults, losses given default etc). These complex, non-linear

channels of risk propagation mean the probability distributions generated from RAMSI are fat-tailed (Chart 17).

These ingredients, with rich interactions both within and across banks and between banks and the economy, mean RAMSI is a useful framework when undertaking top-down stress-testing. Alongside bottom-up granular models, the Bank used RAMSI in its concurrent stress-testing of UK banks at the end of last year. Because it can be used to draw probability distributions of bank balance sheet measures, such as liquidity and solvency, in time this may enable probabilistic “fan charts” of bank strength to be constructed. These would be akin to the fan charts for inflation and output produced in the Bank’s *Inflation Report* for monetary policy purposes.

In a broadly similar spirit, but at a greater level of granularity, Agent-Based Models (ABM) are a natural vehicle for modelling behavioural interactions within and across macro-financial systems (Farmer and Foley (2009)). There has been some progress towards developing ABMs, in particular since the financial crisis. For example, the CRISIS (Complexity Research Initiative for Systemic Instabilities) project is a consortium of EU universities, private firms and policymaking institutions which is building large-scale calibrated models to capture sectoral behaviour, including between the economy and the financial system.¹²

At national level, promising progress has also been made in developing ABM models of the housing market. For example, Geanakoplos et al (2012) develop a regional model of the US housing market, which does a much better job than aggregate macro-economic models in matching housing market patterns, both in the run-up to, and following, the financial crisis.

As part of its new research agenda, the Bank of England is working currently with the Institute for New Economic Thinking (INET) at Oxford to try and develop an ABM model of the UK housing market.¹³ As with Geanakoplos et al (2012), the aim is to provide a calibrated, granular model which better captures the dynamics of the housing market. It could also, potentially, be used to help evaluate the efficacy of macro-prudential interventions in that market – for example, policies which alter mortgage loan-to-income or loan-to-value limits.

A third promising area of research concerns the interaction the macro-prudential and macro-economic layers of the system – that is, the interplay between macro-prudential and monetary policy tools. This is another priority area for the Bank’s new research agenda. Although these policies have distinct primary objectives, their transmission channels are likely to be closely interwoven, with both affecting risk-taking and economic activity. This adds additional complexity to the setting of monetary and macro-prudential policies.

It is possible to explore this complex interplay using simulations from models comprising both nominal and financial frictions. Both are needed to accommodate a role for monetary policy (whose effectiveness relies primarily on nominal frictions) and macro-prudential policies (whose effectiveness relies principally on financial frictions). A number of models have recently been developed with these core ingredients and have been used to explore monetary and macro-prudential policy interactions (Smets (2013), Paustian and De Paoli (2013)).

Chart 18 shows the results of a simulation from a model developed by two colleagues at the Bank (Aikman and Nelson (2014)). This, too, embodies both nominal and financial frictions. On the axes are plotted two measures of stability – macro-economic stability (as proxied by the variability of nominal GDP) on the x-axis, financial stability (as proxied by the variance of

¹² More details are available at www.crisis-economics.eu.

¹³ For more on the Bank’s research agenda, please see <http://www.bankofengland.co.uk/research/Documents/onebank/discussion.pdf>.

credit spreads) on the y-axis. Historically estimated levels of these two variables, normalised to one, are shown at point A in Chart 18.

If we now allow monetary policy to play an active role in stabilising the business cycle, through an optimal policy rule which weights output and inflation deviations from target, we can move the economy to point B. This is unambiguously preferred to point A, with greater stability in both the financial system and, in particular, the wider economy. But could we do better still by having monetary policy assume *explicit* responsibility for safeguarding financial stability?

If we augment this monetary policy rule to take account of financial stability factors, by having interest rates respond to credit spreads, we move to point C. This has the benefit of improving financial stability relative to the conventional monetary policy case. But this comes at the cost of destabilising somewhat the macro-economy. It is unclear whether point C dominates B in a welfare sense. In other words, having monetary policy meet both macro-economic and macro-prudential objectives involves a trade-off.

Once we add an explicit macro-prudential instrument to the equation, however, this trade-off disappears, or at least is lessened. In particular, let's add a regulatory-set counter-cyclical capital buffer for banks, the like of which is now part of Basel III. Used in tandem with optimal monetary policy, this moves the economy to point D. This involves both greater macro-economic and financial stability than the alternatives. On the face of it at least, it is welfare-improving.

Although the model used is specific, the policy lesson appears to be a general one. Having two instruments (monetary and macro-prudential) leads to an improvement in macro-financial stability. Two policy hands beat one when there are two policy objectives. This in a sense underscores the importance of Tinbergen's augmented policy rule – that there should be as many policy tools as there are complex sub-systems – when managing a complex system of systems.

None of the models so far considers global spill-overs, either through trade or financial channels. Generally-speaking, there is a dearth of macro-economic models which take seriously these international spillovers in an analytically-coherent fashion. Even if modelling this behaviour is over-ambitious in the short-term, there may be merit in improved monitoring of these global flows of financing over time, given their importance for the stability of economies and financial systems.

Elsewhere, I have likened this to the creation of “global weather map”. This could be used to track and map the source, scale and nature of cross-border capital flows in close to real-time (Haldane (2012, 2014) and IMF (2013)). Plotting this complex, rapidly-adapting web would be a natural precursor to using it to address “what if” questions. What, for example, would be the impact of a rise in US interest rates on the international flow of funds? The IMF would be the natural guardians of this surveillance-cum-stress-testing machine.

(c) Policy design

Significant changes have been made to the macro-financial policy architecture since the crisis. In general, these frameworks have moved closer towards meeting Tinbergen's rule of having as many distinct policy tools as there are complex layers in the macro-financial system. These frameworks are now equipped with tools for assessing and addressing risks at a range of resolutions – microscopic, macroscopic and telescopic. And the gaping hole in the pre-crisis policy architecture – the macro-prudential layer – has now been filled.

The institutional architecture for macro-financial policy has also been adapted since the crisis to better enable interactions between these policy layers to be addressed. In the UK, the micro-prudential, macro-prudential and monetary policy arms of policy are now attached to a single body, the Bank of England. Each arm has a separate policy committee tasked with

setting policy. But these policy committees have shared inputs of information and meet jointly to discuss areas of intersecting policy responsibility (Bank of England (2014)).

Taking each of the arms of policy in turn, at the highest level of resolution *micro-prudential policy frameworks* have been significantly revised in a number of countries since the crisis. Often, this has had the aim of refocussing the microscope on the systemic, as well as firm-specific, risks facing banks. For the world's largest banks, significant extra cushions of loss-absorbing capacity are now being required, together with changes to banks' organisational structures to make them credibly resolvable. In some countries, including the UK and US, large banks' activities are being separated or ring-fenced. These initiatives ought, over time, to lower the degree of complexity and inter-connectivity of these entities, making them less virulent super-spreaders of financial contagion.

A second area of progress is stress-testing. A major step forward was taken last year in the concurrent stress-testing of banks across Europe. This followed the successful practice first adopted in the US in 2009. Concurrent stress-testing of the whole banking sector allows both top-down and bottom-up perspectives to be brought to bear when assessing individual firms' risks.

This opens up the possibility of taking seriously complex behavioural interactions, both within the financial sector and between it and the wider economy. For example, a capital-deficient firm might withdraw funding and cause wider liquidity pressures for other banks. Or a sequence of capital-deficient banks might contract lending in a credit crunch, thereby damaging the wider economy and feeding back negatively to other banks' portfolios (Haldane (2009)).

Yet it remains early days in the development of stress-testing technologies. A system of systems perspective could be useful in enabling stress-testing technologies to better capture system-wide risks. For example, interaction and feedback channels could usefully be built into future stress-testing frameworks used by central banks, including the Bank of England, the Federal Reserve and the ECB.

At the next level of policy resolution, a key area of progress in the post-crisis period has been in developing the *macro-prudential framework* and the instruments necessary to execute it (IMF (2013)). Many countries internationally now have a macro-prudential framework in place. And international case law on the efficacy of different macro-prudential tools is being built rapidly (IMF (2013)).

Nonetheless, there is much further to go in understanding and calibrating the impact of macro-prudential tools. As well as country case studies, this is likely to draw on models which explore how micro-economic interactions shape macro-economic phenomena. For example, the Bank has recently begun using a Product Sales Database (PSD) compiled by the UK's Financial Conduct Authority (FCA). This covers almost all UK mortgage transaction since 2006, containing around 13 million transactions.

This large micro database can be used both to understand the dynamics of the UK housing market and to help calibrate macro-prudential interventions to shape it. For example, Chart 19 shows the regional distribution of loans made at a loan-to-income multiple in excess of 4.5 at two dates – 2009Q1 and 2014Q2. Cooler colours signify a smaller proportion of loans. As can be seen, there was a discernible warming in the UK housing market up to 2014, at least for higher loan-to-income loans.

At around this time, the Bank of England's Financial Policy Committee (FPC) made a macro-prudential intervention, constraining to 15% the proportion of mortgages with a loan-to-income multiple in excess of 4.5. This was the first time such a macro-prudential measure had been taken in the UK, though a number of other countries have capped mortgage loan-to-income or loan-to-value multiples (IMF (2013)).

This was not, however, a step into the dark for the FPC. The calibration of its intervention drew explicitly on the micro-evidence contained in the PSD database. This gave a very clear

idea of how many households and banks might be affected by such an intervention and by how much. This was a macro-prudential policy calibrated from micro-prudential constituents. Although a first step, there is clearly further to go in using micro-level data to calibrate the impact of macro-prudential tools on lender and borrower behaviour.

Turning to the *global economic and financial system*, it is here where the existing policy architecture may at present be most deficient. Some have gone further and argued that there is nothing at present that much resembles a global financial architecture at all (de Larosiere (2014), Haldane (2014)). Despite the crisis being the first truly global one, reform of the global financial architecture has been slow.

There are many dimensions to reform of the global financial architecture. These include the appropriate role and resourcing of the IMF, the reserve currency role of the dollar and the appropriate role of so-called for capital flow management policies to modulate fluctuations in the global flow of funds (IMF (2012)). Each of these has been widely discussed and debated over many decades. On some issues, there has been progress. For example, the IMF now endorses the use of capital flow management policies in some situations, in contrast to its position a few years ago.

But there is one area where progress in strengthening the international financial architecture has been much slower. This is how best to deal with common shocks to global safe and risky yields. One striking feature of the past few years has been the extremely high correlation among asset prices globally, in particular among advanced economies. This is true of both “safe” rates of return on government assets and “risky” rates of return on private assets (Chart 20). In either case, correlations are extremely high, hovering around 0.9.

This begs a number of questions, both research and policy. What is the root cause of these correlations? One possibility is portfolio shifts by global asset managers, allocating their portfolio on an asset-by-asset basis. If so, what implications does this carry for national monetary and macro-prudential policymakers, seeking to steer safe and risky rates respectively? Do high global correlations strengthen the case for international co-ordination of monetary and macro-prudential policies?

There may be greater scope to co-ordinate macro-prudential tools. One way of doing so is to develop macro-prudential instruments which operate on an asset-class basis, rather than on a national basis. This would be recognition that asset characteristics, rather than national characteristics, may be the key determinant of portfolio choices and asset price movements, perhaps reflecting the rising role of global asset managers.

There has already been some international progress towards developing asset market specific macro-prudential tools, specifically in the context of securities financing transactions where minimum collateral requirements have been agreed internationally (FSB (2014)). But there may be scope to widen and deepen the set of financial instruments covered by prudential requirements, to give a richer array of internationally-oriented macro-prudential tools. These would then be better able to lean against global fluctuations in a wider set of asset markets.

Conclusion

This time was different: never before has the world suffered a genuinely *global* financial crisis, with every country on the planet falling off the same cliff-edge at the same time. This fall laid bare the inadequacy of our pre-crisis understanding of the complexities of the financial system and its interaction with the wider economy, locally but in particular globally. It demonstrated why the global macro-financial network is not just a complex adaptive system, but a complex system of systems.

The crisis also revealed gaps and inadequacies in our existing policy frameworks. Many of those gaps have since been filled. Micro-prudential microscopes have had their lens refocused. Macro-prudential microscopes have been (re)invented. And global telescopes

have been strengthened and lengthened. Institutional arrangements have also been adapted, better enabling co-ordination between the micro, macro and global arms of policy. So far, so good.

Clearly, however, this remains unfinished business. The data necessary to understand and model a macro-financial system of systems is still patchy. The models necessary to make behavioural sense of these complexities remain fledgling. And the policy frameworks necessary to defuse these evolving risks are still embryonic. More will need to be done – both research and policy-wise – to prevent next time being the same.

References

Aikman, D, Alessandri, P, Eklund, B, Gai, P, Kapadia, S, Martin, E, Mora, N, Sterne, G and Willison, M (2009), “Funding liquidity risk in a quantitative model of systemic stability”, *Bank of England Working Paper No. 372*.

Aikman, D, Haldane, A and Nelson, B (2014), “In search of macroprudential policy rules for the UK”, mimeo.

Aikman, D, Haldane A and Nelson, B (2014), “Curbing the credit cycle”, *The Economic Journal*.

Albert, R, Jeong, H & Barabási, A (2000), “Error and attack tolerance of complex networks”, *Nature* **406**, 378–382.

Arthur, W. Brian (2014), “Complexity and the Economy”, *Oxford University Press*.

Barabási, Albert-László; Albert, Réka. (1999), “Emergence of scaling in random networks”. *Science* **286** (5439): 509–512.

BCBS (2014), “The G-SIB assessment methodology – score calculation” available at <http://www.bis.org/bcbs/publ/d296.pdf>.

Bems, R., Johnson, R.C. and Yi, K. (2012), *The Great Trade Collapse*, NBER working paper 18632.

BIS (2014), Annual Report, No. 84, available at http://www.bis.org/publ/arpdf/ar2014_ec.pdf.

Bisin, A, Moro, A and Topa, G (2011), The Empirical Content of Models with Multiple Equilibria in Economies with Social Interactions, *Federal Reserve Bank of New York Staff Reports*, no. 504.

Burrows, O, Learmonth, D and McKeown, J (2012), “RAMSI: a top-down stress-testing model”, *Bank of England Financial Stability Paper No. 17*.

Carmassi, J & Herring, R (2014) “Corporate structures, transparency and resolvability of global systemically important banks”, The Systemic Risk Council research paper.

Collinson, S, Jay, M (2012), “From Complexity to Simplicity”, Palgrave Macmillan.

Craig, B, Alter, A and Raupach, P (2014), “Centrality-based Capital Allocations”, IMF Working Paper, December 2014.

Davies, R & Tracey, B (2014), “Too Big to Be Efficient? The Impact of Implicit Subsidies on Estimates of Scale Economies for Banks” *Journal of Money, Credit and Banking*.

De Paoli and Paustian (2013), “Coordinating Monetary and Macroprudential Policies”, *Federal Reserve Bank of New York Staff Reports*, No. 653.

Farmer, J.D., and Foley, D (2009), “The Economy Needs Agent-Based Modelling” *Nature* vol. 460.

FSB (2012), “Enhanced Disclosure Task Force”, available at

http://www.financialstabilityboard.org/wp-content/uploads/pr_120510.pdf.

FSB (2014), Strengthening Oversight and Regulation of Shadow Banking Report, available at http://www.financialstabilityboard.org/wp-content/uploads/r_141013a.pdf.

Gai, P, Haldane, A and Kapadia, S “Complexity, Concentration and Contagion”, *Journal of Monetary Economics*, Vol. 58, (5), pages 453–470, 2011.

Gai, P and Kapadia, S (2011), “Contagion in financial Markets”, *Bank of England Working Paper Number 383*.

Gabaix, X (2009), “Power laws in economics and finance”, *Annual Review of Economics*, 1 pp. 255–93.

Gao, J Liu, Y, D’Souza, R & Barabási, A (2014), Target control of complex networks, *Nature Communications* 5, 5415.

Geanakoplos, J, Axtell, R, Farmer, J.D., Howitt, P, Conlee, B, Goldstein, J, Hendrey, M, Palmer and N.M., Yang, C–Y (2012), “Getting at systemic risk via an agent-based model of the housing market”, *American Economic Review* vol.102.

Goldin, I. and Mariathasan, M. (2014), “The Butterfly Defect: How Globalization Creates Systemic Risks, and What to Do about It”, *Princeton University Press, July*.

Gorod, A, White, B, Ireland, V, Gandhi, S and Sauser, B (2014), “Case Studies in System of Systems, Enterprise Systems, and Complex Systems Engineering”, *CRC Press*, Pages 3–32.

Haldane, A. G (2009), “Rethinking the financial network”, available at <http://www.bankofengland.co.uk/publications/speeches/2009/speech386.pdf>

Haldane, A. G (2009), “Banking on the State”, available at <http://www.bankofengland.co.uk/archive/Documents/historicpubs/speeches/2009/speech409.pdf>.

Haldane, A and May, R (2011), “Systemic risk in banking ecosystems”, *Nature*, issue 469, 351–355.

Haldane, A (2012), “Tails of the unexpected”, speech at “The Credit Crisis Five Years On: Unpacking the Crisis”, conference held at the University of Edinburgh Business School, 8–9 June, available at <http://www.bankofengland.co.uk/publications/Documents/speeches/2012/speech582.pdf>.

Haldane, A (2014), “Twin Peaks”, speech at the Kenilworth Chamber of Trade Business Breakfast, 17 October 2014, available at <http://www.bankofengland.co.uk/publications/Documents/speeches/2014/speech716.pdf>.

Hanson, S.G, Kashyap, Anil and Stein, Jeremy (2010), “A Macroprudential Approach to Financial regulation”, Chicago Booth Research Paper, No. 10–29.

Hills, B and Hoggarth, G (2013), Cross border bank credit and global financial stability, Bank of England Quarterly Bulletin Q2, 2013.

IMF (2013), “Key Aspects of Macroprudential Policy”, available at <http://www.imf.org/external/np/pp/eng/2013/061013b.pdf>.

Jackson, M O (2010), “Social and economic networks”, Princeton.

Kurant, M and Thiran, P (2006), Layered Complex Networks, *Physical Review Letters* 96.

Ladyman, J, Lambert, J & Wiesner, K (2013), “What is a Complex System?”. *European Journal for Philosophy of Science*, vol 3., pp. 33–67.

Langfield, S, Liu, Z and Ota, T (2014), “Mapping the UK interbank system”, Bank of England Working Paper, No. 516.

- Liu, Y, Slotine, J and Barabási, A.L (2011)**, Controllability of complex networks, *Nature* 473, 167–173 (2011) .
- Minoiu, C. and Reyes, J.A. (2013)**, A network analysis of global banking: 1978–2010, *Journal of Financial Stability* 9 (2013) 168–184.
- Newman, M E J (2005)**, “Power laws, Pareto distributions and Zip’s law”, *Contemporary Physics*, 46 (5), pp. 323–51.
- Newman, M. E. J. (2011)**, “Complex Systems: A Survey”, *Am. J. Phys.* 79, 800–810.
- OECD (2013)**, *Interconnected Economies: Benefiting from Global Value Chains*, OECD.
- Perrow, C (1984)**, *Normal Accidents: Living with High-Risk Technologies*, Basic Books, NY.
- Pozsar, Z, Adrian, T, Ashcraft, A and Boesky, H (2010)**, “Shadow Banking”, *Federal Reserve Bank of New York Staff Reports*, No. 458
- Roger, S and Vicek, J (2011)**, Macrofinancial Modeling at Central Banks: Recent Developments and Future Directions, *IMF Working Paper* 12/21.
- Shakir, T and Tong, M (2014)** The interaction of the FPC and the MPC, *Bank of England Quarterly Bulletin* Q4, available at <http://www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2014/qb14q403.pdf>.
- Simon, H (1962)**, The Architecture of Complexity, *Proceedings of the American Philosophical Society*, Vol. 106, No. 6., pp. 467–482.
- Simplicity consulting (2012)**, “True Cost of complexity in the banking sector” available at http://www.simplicity-consulting.com/wp-content/files_mf/1340875856BankingTrueCostofComplexity2012.pdf.
- Smets, F (2013)**, “Financial Stability and Monetary Policy: How Closely Interlinked?” *Sveriges Riksbank Economic Review*.