



Munich Personal RePEc Archive

Taking Uncertainty Seriously: Simplicity versus Complexity in Financial Regulation

David Aikman and Mirta Galesic and Gerd Gigerenzer and
Sujit Kapadia and Konstantinos Katsikopolous and Amit
Kothiyal and Emma Murphy and Tobias Neumann

Bank of England, Max Planck Institute for Human Development

2 May 2014

Online at <https://mpra.ub.uni-muenchen.de/59908/>

MPRA Paper No. 59908, posted 14 November 2014 20:31 UTC

TAKING UNCERTAINTY SERIOUSLY:
SIMPLICITY VERSUS COMPLEXITY IN FINANCIAL REGULATION¹

By

David Aikman, Mirta Galesic, Gerd Gigerenzer, Sujit Kapadia, Konstantinos Katsikopoulos,
Amit Kothiyal, Emma Murphy and Tobias Neumann

2 May 2014

Abstract: Distinguishing between risk and uncertainty, this paper draws on the psychological literature on heuristics to consider whether and when simpler approaches may out-perform more complex methods for modelling and regulating the financial system. We find that: (i) simple methods can sometimes dominate more complex modelling approaches for calculating banks' capital requirements, especially if limited data are available for estimating models or the underlying risks are characterised by fat-tailed distributions; (ii) simple indicators often outperformed more complex metrics in predicting individual bank failure during the global financial crisis; (iii) when combining information from different indicators to predict bank failure, "fast-and-frugal" decision trees can perform comparably to standard, but more information-intensive, regression techniques, while being simpler and easier to communicate.

¹ All authors: David Aikman, Bank of England, Threadneedle Street, London, EC2R8AH, david.aikman@bankofengland.co.uk; Mirta Galesic, Max Planck Institute for Human Development, Berlin, galesic@mpib-berlin.mpg.de; Gerd Gigerenzer, Max Planck Institute for Human Development, Berlin, gigerenzer@mpib-berlin.mpg.de; Sujit Kapadia, Bank of England, Threadneedle Street, London, EC2R8AH, sujit.kapadia@bankofengland.co.uk; Konstantinos Katsikopoulos, Max Planck Institute for Human Development, Berlin, katsikop@mpib-berlin.mpg.de; Amit Kothiyal, Max Planck Institute for Human Development, Berlin, amit.kothiyal2004@gmail.com; Emma Murphy, Bank of England, Threadneedle Street, London, EC2R8AH, emma.murphy@bankofengland.co.uk; Tobias Neumann, Bank of England, Threadneedle Street, London, EC2R8AH, tobias.neumann@bankofengland.co.uk.

We would like to thank Mervyn King for instigating this research collaboration. We are grateful to Ryan Banerjee, Bob Chirinko, Renzo Corrias, Jas Ellis, Andy Haldane, Simon Hall, Ramesh Kapadia, Vasileios Madouros, Hitoshi Mio, Marco Raberto, Tarik Roukny, Vicky Saporta, Jean Whitmore and seminar participants at the Bank of England, the London School of Economics, the Deutsche Bundesbank / SAFE conference on 'Supervising banks in complex financial systems' (Frankfurt, 21-22 October 2013) and the ESRC conference on 'Diversity in macroeconomics' (Colchester, 24-25 February 2014) for helpful conversations and comments. We would also like to thank George Avil, Timothy Richards, Gowsikan Shugumaran, George Slade, and Charlie Woolnough for excellent research assistance.

Disclaimer: The views expressed in this paper are those of the authors, and are not necessarily those of the Bank of England or Financial Policy Committee members.

“It is better to be roughly right than precisely wrong”

– John Maynard Keynes (attributed)

1 Introduction

The financial system has become increasingly complex over recent years. Both the private sector and public authorities have tended to meet this complexity head-on, whether through increasingly complex modelling and risk management strategies or ever-lengthening regulatory rulebooks. But this neither helped to predict, nor to prevent, the global financial crisis.

The dominant paradigm for studying decision-making in economics and finance is based on rational agents who operate in situations of known, calculable risks. There is no cost to complexity in such a setting: more information is always perceived to be better than less; and decisions should optimally weight all relevant factors. The result has been a quest for ever-greater precision – and hence ever increasing complexity – in the models and toolkits typically being developed and used in applied work. This is also reflected in elements of the approach towards banking regulation which allows banks to use their own internal models to calculate regulatory capital requirements based upon underlying estimates of variables such as default probabilities and losses in the event of default, and has led to an exponential rise in the number of calculations required for a large, universal bank from single figures a generation ago to hundreds of thousands, perhaps even millions, today.

But many real-world problems do not fall neatly into the category of known, calculable risks that such approaches are designed for. The likelihood of a systemic financial crisis occurring over the next year, for instance, involves so many unpredictable factors as to be unknowable. Problems such as these are better characterised by ‘Knightian’ uncertainty, rather than risk, a distinction Frank Knight (1921, p. 233) established in his seminal work *Risk, Uncertainty and Profit*:

“The practical difference between the two categories, risk and uncertainty, is that in the former the distribution of the outcome in a group of instances is known, ... while in the case of uncertainty this is not true...”

Knight’s examples of risks include cases in which probabilities can be estimated *a priori*, such as the throwing of dice, or by sampling, such as calculating the risk of bottles bursting in a champagne factory. By contrast, uncertainty arises when not all risks are known or even knowable. The assumption that decision makers follow consistent, rational rules becomes untenable in a world of uncertainty. This raises the question of whether we need a rather different set of tools to deal with such problems?

The central premise of this paper is that the distinction between risk and uncertainty is crucial and has received far too little attention from the economics and finance professions to date.² The shift from risk to uncertainty can turn what we think we know about decision making upside down. Decision rules that attempt to achieve ever greater precision can become increasingly *imprecise*; rules that attempt to weight optimally all the relevant information can sometimes generate poorer results than those based on simple averages or those that deliberately disregard information. Taking uncertainty seriously forces us to recognise that, in some circumstances, there are potential benefits to more simplicity over greater complexity.

To explore these issues further, this paper draws on lessons from the psychological literature on heuristics (Gigerenzer and Brighton, 2009) and considers how a heuristic approach may be a complementary tool for dealing with uncertainty in financial regulation. In particular, the paper argues that, in the face of uncertainty, adding complexity may sometimes lead to poorer

² Notable exceptions include Caballero and Pindyck (1996) , Hansen and Sargent (2007), and in the context of the financial system, Haldane and Madouros (2012).

performance. Three main findings are reported to support this claim. First, simple methods can sometimes dominate more complex modelling approaches for calculating banks' capital requirements. According to simulation results, this is more likely to be the case when limited data are available for estimating models and the underlying risks are characterised by fat-tailed distributions. Second, on an individual basis, simple indicators, such as leverage or loan-to-deposit ratios, often outperformed more complex metrics in predicting failure across a cross-country sample of large banks during the global financial crisis. And third, when combining information from different indicators to predict bank failure, "fast-and-frugal" decision trees, which deliver a simple classification scheme, can perform comparably to standard, but more information-intensive, regression techniques, while being simpler and easier to communicate.

The rest of the paper is organised as follows. Section 2 discusses the economics of risk and uncertainty and the tools available for handling both environments. Section 3 explains the heuristics approach and sets out what we know about the conditions under which they can provide an effective tool for decision making. In Section 4, these ideas are applied to a specific financial policy question: how to measure risk for the purposes of calculating banks' capital requirements. In Section 5, the focus is on the empirical prediction of bank failure during the global financial crisis, illustrating how the use of simple indicators and approaches may sometimes lead to superior performance. Section 6 concludes.

2 The economics of risk and uncertainty

2.1 A world of risk

The question of how best to model choice under uncertainty lies at the core of economic and finance theory. But the two bedrocks of modern microeconomics and finance model the unknown future by taking a risk approach – assuming full knowledge about outcomes and their associated probability. In the general equilibrium framework of Arrow and Debreu (1954), agents are assumed to be rational and know the probability distribution of all future states of the world. Any predictions about behaviour are free from psychological or sociological factors and behaviour should always converge to the ideal choice predicted by rational choice theory. Risk can be perfectly assessed and thereby priced, traded and hedged. Markowitz (1952) and Merton (1969) also assume a known probability distribution for future market risk. This enables portfolio risk to be calculated exactly.

Together these frameworks are purported to help explain patterns of behaviour from consumption and investment to asset pricing and portfolio allocation.

The future is rarely so straightforward. Yet in finance, the dominant portfolio allocation and pricing models remain built on mean-variance foundations and the principles of quantifiable risk. They are also often firmly rooted in normality (Black and Scholes, 1973): the normal distribution provides an appealing, simple description of the world with outcomes lying in a perfect bell-shape symmetrically around a mean. But this assumption of normality can result in a massive underpricing of catastrophe risk. As a result, by assuming normality, traders relying on Black-Scholes systematically misprice options, believing they are cheaper than some intrinsic value. Given this, such options are sold in greater size than they ought to be. Since tail events, by definition, happen infrequently, this is rarely an issue most of the time. But when such events do crystallise, as demonstrated since 2007, the ramifications of ignoring uncertainty can make for dramatic adjustments in price.

Of course, the fact that true distributions can be more complex than most standard models assume is not new. When faced with evidence that the basic assumptions and predictions of models were flawed, the financial industry responded with even more complex models that attempted to ‘patch’ the old ones rather than acknowledge the shortcomings (Mandelbrot and Hudson, 2004). For example, what is unlikely in, say, a normal distribution might not be unheard of in a distribution with fatter tails – black swans might simply be more common than expected. In these cases, all that is left to do is to find the right distribution. Adding complexity, it might be argued, is then a good thing as it makes the models more ‘realistic’.

But such adjustments lead to ever more complex models that are impenetrable for most managers who nevertheless need to use them to make everyday financial decisions. This is compounded by the fact that many of the ‘quants’ who designed the models often lacked the practical experience of financial markets that would have helped them to understand how different the markets are from their models (Derman and Wilmott, 2009).

Even such complex, adjusted models can break down once uncertainty is introduced into the system. In fact, the use of an overly flexible model can itself contribute to instability. The problem with many financial models is not merely that they attempt to describe a phenomenon, as many other theories in natural science do, but they can also become the bedrock theory on which financial

engineering and decision making is rooted. As a result, the models enter into the functioning of the system being described, and so have the potential to destabilise markets and exacerbate financial risks – that is, they become part and parcel of the underlying problem itself (Lucas, 1976; Caccioli, Marsili and Vivo, 2009).

2.2 *Adjusting for uncertainty*

What if a distribution is simply unknowable – either intrinsically or because of practical limitations?

There are several reasons why the behaviour of financial systems might be characterised by uncertainty (Aikman *et al*, 2011). First, assigning probabilities is particularly difficult for rare, high-impact events, such as financial crises, because there are few precedents and the causal mechanisms are not well understood. This means that as understanding of these processes develops, the assessment of their likelihood may change, possibly sharply.

Second, the behaviour of financial systems can also be very sensitive to small changes in initial conditions and shocks, which may lead to very different outcomes. This could be because they exhibit chaotic dynamics or are subject to multiple equilibria, which can lead to strong path-dependency or hysteresis. But it also reflects the central role of network and feedback effects in propagating financial contagion. Complex systems can exhibit ‘tipping points’, where for a small change in parameter values, the system can move from a state in which contagion dies out to one in which it spreads through the entire population. Such results are widely appreciated in epidemiology (Anderson and May, 1991) but recent analysis shows how they also apply to financial systems (Gai and Kapadia, 2010; May and Arinaminpathy, 2010; Gai, Haldane and Kapadia, 2011). Moreover, in such setups, *a priori* indistinguishable shocks – such as the failure of two different, but identical-looking, banks – can have vastly different effects on system stability depending on the position in the network of the banks concerned or the state of other members of the network. So the interconnected nature of modern globalised finance can mean that financial risk may pass quickly and extensively through the system in unpredictable ways.

Third, in stark contrast to complex physical systems, economic and financial systems may be highly unpredictable because they involve human actors whose beliefs about the past, present and future shape their behaviour and thus economic outcomes (see also Subbarao, 2011). Key variables that drive financial systems thus reside in people’s minds and are, therefore, necessarily unknowable. For example, if financial market participants are uncertain about the state of the economy, they may

be more conservative in their risk-taking; this may reduce growth and weaken the economy further. Similarly, in a bull market, irrational exuberance can encourage a feeling that the good times will never end. Such beliefs adapt over time in response to changes in the environment. As a result, there may be few, if any, genuinely ‘deep’ or ‘stable’ underlying parameters or relationships in economics and finance, with no model being able to meet the Lucas (1976) critique.

These factors mean there could be no way in which we could even quantify the probability of meeting a ‘black swan’, ie a large impact, unforeseen, random event (Taleb, 2007). And if we cannot determine a probability distribution, we are confronted not with risk, but with ‘Knightian’ uncertainty.

2.3 *Tools for managing uncertainty*

Many scientific disciplines are confronted by uncertainty and have begun to develop tools to deal with the issue. An emerging solution in other disciplines is to build in robust, simple strategies that can handle such unpredictability. For example, technological products which are designed to function well under a broad range of user conditions are likely to be more profitable in a business context (Taguchi and Clausing, 1990). In biology, animals often use amazingly simple and efficient approaches to solve complex problems such as finding a mate or a nesting location – peahens choose their mate by investigating only three or four peacocks in a large lek, and choose the one with the largest number of eyespots (Petrie and Halliday, 1994). And ants estimate the area of a potential nest cavity by running around it and leaving a pheromone trail, and after a while running around it again but on a different path. The size of the cavity is proportional to the frequency of encountering the old trail (Mugford, Mallon and Franks, 2001).

In forecasting, Makridakis and Hibon (1979, 2000) have shown that a simple time series model sometimes out-predicts many complex and statistically sophisticated models that use many more variables. Their results suggest that while complex models can fit the data well, their predictive power is sometimes poor. Fitting corresponds to known risks, whilst prediction involves an element of uncertainty.

Behavioural economics has also tried to overcome some of the issues presented by uncertainty. Experimental evidence shows persistent deviations in the decision-making and behaviour exhibited by individuals and firms from the assumptions and predictions of neoclassical theory (Camerer,

1999; Kahneman and Tversky, 1979; Rabin, 1998). But the literature has responded to these experimental findings in different ways. One strand of the literature has sought to preserve expected utility theory and other key neoclassical assumptions as the normative standards for evaluating human decision making. Deviations from the neoclassical paradigm are viewed by researchers in this strand as sub-optimal or irrational in some broad sense (Kahneman, 2011).

A second, closely related, strand attempts to encapsulate the lessons from the experimental evidence into simple parameters or theories which are then incorporated back into mainstream neoclassical models. For example, Kahneman and Tversky's prospect theory (1979) attempts to 'repair' the expected utility paradigm by including additional parameters and transformations of probabilities and outcomes. While there is value in this approach – recognising the genuine costs of collecting information or the psychological underpinnings of non-standard preferences, for example – the theories typically apply to choices made under risk, not uncertainty. Moreover, ignoring or simplifying information cannot be 'optimal' from this perspective (Gigerenzer and Selten, 2001).

A third, rather different strand builds on the insights of Simon (1955), who believed that human behaviour followed simple rules precisely because humans operate in complex environments. Simon provided an important link between psychological reality and decision-making by introducing the concept of 'bounded rationality' to explain how people seek satisfaction, instead of maximizing utility, as conventional economics presumed. Simon's bounded rationality is neither optimisation (under constraints) nor irrationality – minds with limited time, knowledge and other resources can still attain successful outcomes by exploiting features of their environments to form heuristics. A body of literature in economics builds on this line of thinking by developing simple models using heuristics that incorporate a role for beliefs to explain phenomena such as multiple equilibria and path dependence (eg Bikhchandani *et al*, 1992).

3 What are heuristics and when do they work?

Humans and other animals rely on heuristics to deal with an uncertain world (Gigerenzer, Hertwig and Pachur, 2011). In a human context, simple heuristics have been applied to decision making in diverse fields such as management, medicine and engineering (see Katsikopoulos, 2011, for a review). A heuristic is a simple rule that ignores part of the available information, in order to make inferences, which turn out to be more robust and accurate in many cases. An example of such a rule – a fast-and-frugal tree – is described in Box 1. In a similar vein, lexicographic models such as

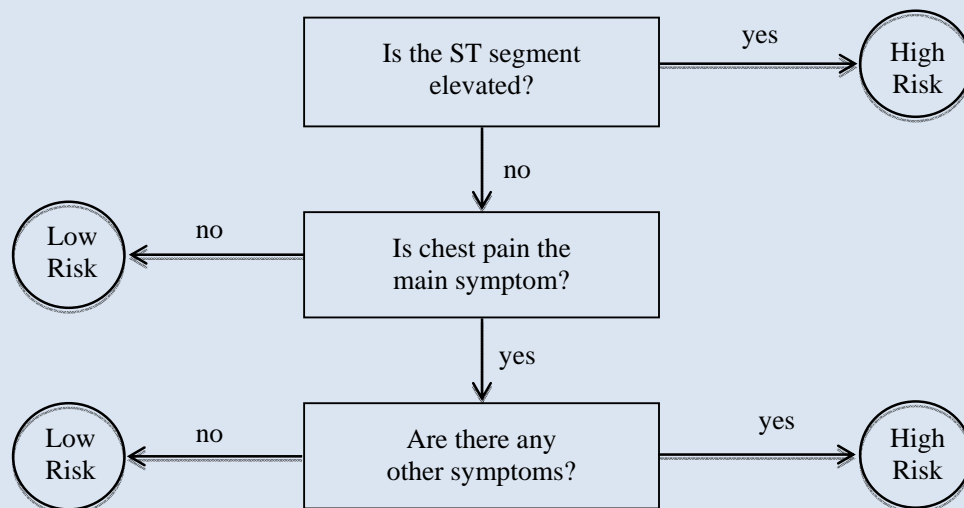
‘take-the-best’ (a heuristic that chooses one of two options based on a single cue, such as people choosing a restaurant that is full over one that is empty despite similar menus) have often been found to be superior in predictive accuracy, compared to both regression (Czerlinski, Gigerenzer and Goldstein, 1999) and Bayesian models (Martignon and Hoffrage, 2002). This is particularly so when datasets available to fit models are small (Katsikopoulos, Schooler and Hertwig, 2010). And the ‘hiatus heuristic’, which predicts the likelihood that a customer will make a purchase based upon how recent their last purchase was, has been found to outperform more complex models in prediction (Wübben and Wangenheim, 2008).

Heuristics are not pathological deviations from axiomatic rationality, but rather provide valuable procedures for making a decision well in certain complex circumstances involving uncertainty (see Gigerenzer, Hertwig and Pachur, 2011 for further details on the conceptual foundations of heuristics and their formal implementation). The heuristics we discuss can be expressed by precise models that make precise predictions.

Box 1: Fast-and-frugal trees

Doctors can improve their ability to correctly assign potential heart attack victims between intensive and normal care facilities by using fast-and-frugal trees (Green and Mehr, 1997). This decision tree is shown in Figure A below.

Figure A: A fast-and-frugal tree for categorizing patients as having a high, or low, risk of ischemic heart disease



Source: Based on Green and Mehr (1997).

The tree first asks the doctor to assess whether the patient’s electrocardiogram shows an elevated ‘ST segment’. If the answer is positive, no other information is required and the patient is categorised as having a high risk of heart failure, and is thus assigned to intensive care. If the ST segment is not elevated, the doctor is then asked to assess whether chest pain is the patient’s main symptom or not. This time a negative answer provides the exit — the patient is immediately categorised as low risk, and assigned to normal care. Finally, the doctor is asked to consider if there are any other symptoms from a short additional pre-specified list – if so, the patient is categorised as high risk; if not, she is categorised as low risk. The tree is said to be ‘fast and frugal’ because it uses just a few pieces of information, not all of which are always used (Martignon, Katsikopoulos and Woike, 2008). It can out-perform logistic regression, which exploits all of the information, in assigning patients correctly.

3.1 When might heuristics work?

In a world of known risk (with known probability distributions), there is an accuracy-effort trade off. That is, a heuristic that ignores part of the information cannot make more accurate predictions than a more complex model. But once uncertainty is introduced (with unknowable probability distributions), the accuracy-effort trade off no longer necessarily applies. Here, simple heuristics can sometimes do better than more complex models – less is more.


One framework to understand this effect is the bias-variance trade-off (Brighton and Gigerenzer, 2012; Geman *et al*, 1992; Hansen *et al*, 1953). Prediction error can be decomposed into bias, variance, and noise:

$$\text{prediction error} = (\text{bias})^2 + \text{variance} + \text{noise},$$

where *bias* is the difference between the mean estimated function and the true function describing the data; *variance* is the mean squared difference between individual estimates from different samples and the mean estimated function; and *noise* is irreducible error, such as measurement error.

In general, complex models will have better fit and therefore low bias. But with many free parameters to estimate from small samples, complex models also run the danger of over-fitting the parameters to idiosyncrasies of each individual sample, resulting in larger overall variance across samples. On the other hand, heuristics tend to have larger bias because they ignore information. But with few or no free parameters, they typically have lower variance than more flexible complex models. The variance of the more flexible complex models is often so large for small sample sizes that it overshadows the error of heuristics due to bias. Figure 1, taken from Gigerenzer and Brighton (2009), provides a simple example.

Figure 1: Illustrating the bias-variance trade off



Source: Gigerenzer and Brighton (2009).

The plot on the left shows the (unknown) underlying ‘true’ function of a variable at each day of one year which is a degree-3 polynomial, $h(x)$, along with a random sample of 30 noisy observations of $h(x)$ (all based on actual mean daily temperatures in London in 2000). The plot on the right shows, as a function of the degree of polynomial, the mean error in predicting the entire function after fitting polynomials to samples of 30 noisy observations. This error is decomposed into bias and variance, also plotted as functions of the degree of polynomial.

While the 30 observations would be best fitted with a high degree polynomial, the polynomial of degree 3 achieves the highest predictive accuracy. The more complex polynomials with higher degrees suffer from too much variance. On the other hand, polynomials with degrees 1 or 2 are too simple to predict well; that is, they suffer from too much bias. The increase in error with polynomials of degree 4 and higher illustrates a ‘less-is-more’ effect: up to a point, a simpler model

leads to smaller error in prediction than a more complex model. Strikingly, despite both being ‘incorrect’, a polynomial of degree 2 achieves a lower mean prediction error than a polynomial of degree 10.

One example of how far simple arithmetic can take a decision maker is given by the $1/N$ heuristic, used by many investors to allocate wealth across financial assets (Benartzi and Thaler, 2001). As the name suggests, the $1/N$ heuristic – also known as naïve diversification – suggests allocating an equal amount of wealth to each of the assets in one’s portfolio. This heuristic ignores information on the returns of the assets in the past or mean-variance trade-offs – it has bias, but no variance, because it does not estimate any parameters and is therefore insensitive to peculiarities of samples. Empirical evidence from computer simulations suggests that $1/N$ typically outperforms complex strategies, like Markowitz’s (1952) mean-variance optimisation, unless the sample size is very large. For example, DeMiguel *et al* (2007) find that, for a sample threshold of $N = 25$, complex rules outperform simple ones only for sample sizes of in excess of 3000 months (250 years) of data. Intuitively, while the more complex approach may perform better in specific instances, the simpler approach is more robust to large, unpredictable shifts in the value of certain assets, whether railways in the 19th century or dotcom stocks and subprime mortgages more recently. Despite originating one of the key complex models for portfolio theory, Markowitz himself pursued a $1/N$ strategy when investing for his retirement (Bower, 2011).

Bias-variance trade-offs need to be considered in the design of financial regulation too. The lesson here is that the design of regulatory rules needs to be robust in the sense of avoiding over-fitting and creating less variance, even potentially at the cost of stronger bias, because lower overall prediction errors can yield a more prudent regime. To determine whether to use a particular heuristic tool or a more complex strategy, their performance in different environmental circumstances needs to be studied. In other words, there is a need to learn about the ecological rationality of different strategies. The following sections help us to understand when less is more and when more is more.

4 A case study of risk-based capital requirements

This section analyses the trade-offs between complexity and simplicity in the design of bank capital requirements. The aim is to explore the conditions under which a simple heuristic for setting capital requirements, such as giving all assets the same capital charge regardless of their underlying characteristics – a $1/N$ approach – may outperform or underperform a complex approach that seeks

to calibrate capital closely to risks. Capital requirements provide a good case study because the capital framework itself has undergone a gradual evolution towards greater complexity, as we now discuss.

4.1 *Ever-increasing complexity: the historical evolution of capital requirements*

Minimum capital requirements have been co-ordinated internationally since the first Basel Accord of 1988.³ Under Basel I, a bank's assets were allotted via a simple rule of thumb to one of four broad risk categories, each with a fixed 'risk weighting' that ranged from 0-100%. A portfolio of corporate loans, for instance, received a risk weight of 100%, while retail mortgages – perceived to be safer – received a more favourable weighting of 50%. Minimum capital was then set in proportion to the weighted sum of these assets

$$\text{minimum capital requirement} = 8\% \times \sum \text{risk weighted assets.}$$

Over time, this approach was criticised for being insufficiently granular to capture the cross-sectional distribution of risk. All mortgage loans, for instance, received the same capital requirement without regard to the underlying risk profile of the borrower (eg the loan to value ratio, repayment history etc). This led to concerns that the framework incentivised 'risk shifting': as risk was not being 'priced', it was argued that banks had an incentive to retain only the highest risk exposures on their balance sheets as these were also likely to offer the highest expected return.

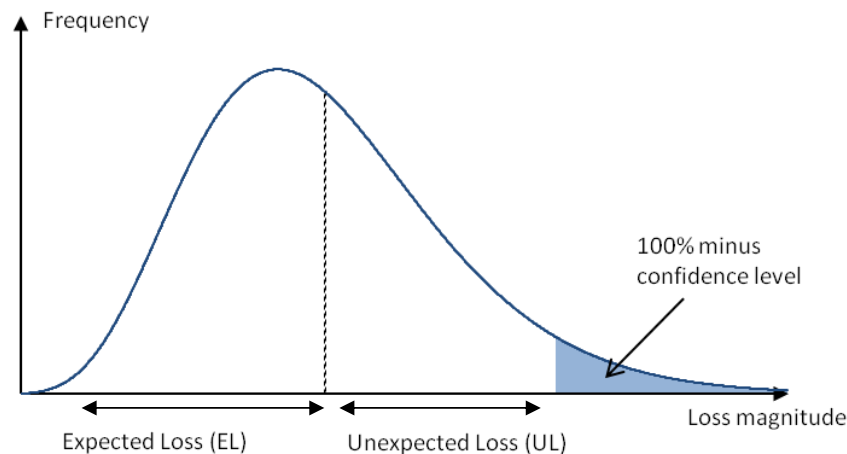
In response, the Basel Committee published a revised set of rules in 2004. The simple four-bucket approach was replaced by one that sought to tie capital requirements much more closely to risks. Banks were encouraged, subject to regulators' approval, to use the Internal Ratings Based approach, under which requirements were based on the outputs of banks' own rating systems under one of two options – a relatively simpler 'Foundation' approach, which is the focus of analysis in this paper, and an 'Advanced' approach, in which banks were allowed to make greater use of internal data to estimate a wider range of parameters. Banks lacking the capacity to model these

³ For a historical perspective on the evolution of capital requirements, see Tarullo (2008).

risks were required to use the Standardised Approach, under which capital requirements were based on external agency ratings or other simple rules of thumb.

The Foundation Internal Ratings Based approach (henceforth simply referred to as IRB) is complex and requires some explanation. Capital requirements are calculated in three steps.⁴ First, banks use their own models to produce an ordinal ranking of borrowers, grouped into a discrete set of rating grades. Second, banks estimate the average probability that borrowers within each grade will default, ie be unable to repay their debts, over the course of the next year. This is called the probability of default or PD.⁵ Third, a given formula sets the capital requirement such that stressed losses will not exceed the bank’s capital up to a 99.9% confidence level. The framework itself contains a back-stop floor that prevents capital requirements from falling below 0.03%. If the assumptions behind the formulae are correct, the output is a capital requirement sufficiently large that a bank is expected to become insolvent only once every 1000 years.⁶ This so-called Value-at-Risk approach is illustrated in Figure 2.

Figure 2: Value-at-Risk model for credit risk under Basel II



The value-at-risk is the loss at the border where the unshaded and shaded area meet (EL + UL). The distribution of the shaded area is not taken into account when calculating value-at-risk.

⁴ We explain the IRB approach using a corporate portfolio – the actual IRB formula in Table 1, floors and precise steps in the capital calculation vary slightly between asset classes. See BCBS (2005) for further explanation of the Basel II IRB formulae.

⁵ Banks on the Advanced IRB approach must also estimate the expected loss given default or LGD of each rating grade – that is, one minus the amount that would be recovered in the event of default.

⁶ In practice though, the output of the formula was scaled up by the Basel Committee on Banking Supervision so that capital requirements of the G10 banking system as a whole were, in aggregate, the same as under Basel I.

Table 1 summarises these distinct approaches in the context of a corporate loan portfolio. They can be ordered on a simplicity-complexity spectrum. Basel I sits at one end of this spectrum: with its single risk weight for all corporate loans, it is the simplest of the three approaches. Using the language of Section 3, it has similarities with the $1/N$ heuristic whereby all loans within an asset class are weighted equally. The Basel II Foundation IRB approach sits at the opposite end of this spectrum and is the most complex approach of the ones we analyse (though, as noted above, the Advanced approach is more complex still). The Basel II Standardised Approach sits somewhere in between: it has five gradations of fixed risk weights mapped against the ratings given by external agencies, as shown in Table 1.

Table 1: Three approaches to calculating capital requirements for corporate loans

Increasing complexity →

	Basel I	Basel II Standardised Approach					Basel II Internal Ratings Based Approach (unexpected loss) ⁷
Risk weight	100%	AAA to AA-	A+ to A-	BBB+ to BB-	Below BB-	Unrated	$K = LGD \left[N \left[(1-R)^{-0.5} \cdot G(PD) + \left(\frac{R}{1-R} \right)^{0.5} \cdot G(0.999) \right] - PD \right]$ $\cdot \left(\frac{1+(M-2.5)b}{1-1.5b} \right)$ $R = 0.12 \cdot \left(\frac{1-e^{-50PD}}{1-e^{-50}} \right) + 0.24 \left(1 - \frac{1-e^{-50PD}}{1-e^{-50}} \right)$ $b = (0.11852 - 0.05478 \ln(PD))^2$
		20%	50%	100%	150%	100%	
Capital requirement, K , on a £100 loan	£8	£1.60	£4	£8	£12	£8	

In what follows, the performance of each of these approaches is compared in an *out-of-sample* forecasting exercise. The purpose is to explore the conditions under which the variance generated by estimating the parameters of the IRB model outweighs the bias of the simpler Basel I and Basel II Standardised approaches. It should be emphasised that the focus is solely on the effectiveness of different approaches towards risk *weighting* under the different systems – in an absolute *levels* sense, it is clear that the Basel I regime in place before the global financial crisis led to a dangerously under-capitalised banking system. Throughout this exercise, we abstract completely from the distinct concern that IRB models may be subject to risk weight ‘optimisation’. In this

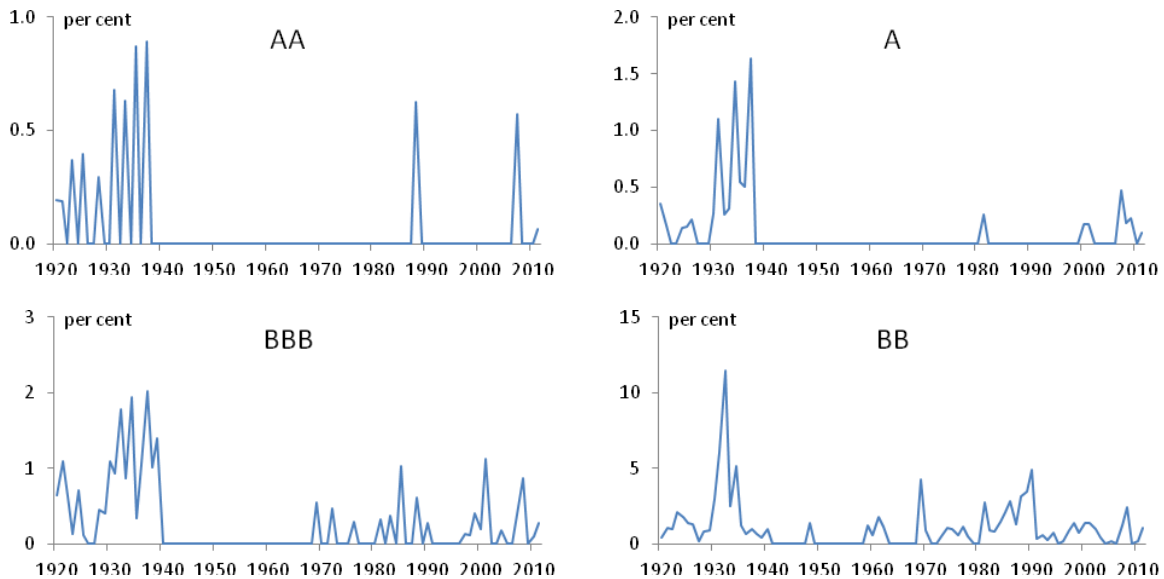
⁷ See BCBS (2006) for an explanation of the parameters used in the formula.

potential scenario, banks might try to identify model-generated parameters partly on the basis that they may yield lower capital requirements.

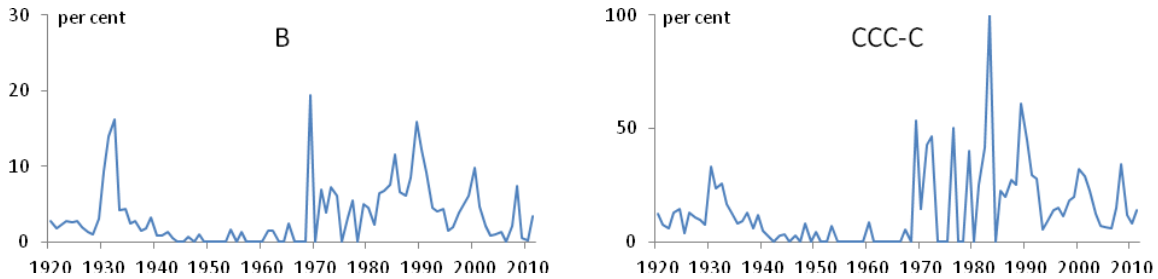
4.2 Data and Methodology

A data generating process is estimated to simulate the occurrence of defaults in a representative corporate loan portfolio. Historical data on bond default rates are used for this purpose, taken from Moody's Investors Service.⁸ The data report annual default rates by rating on senior unsecured corporate bond issues (including financial and non-financial issuers) between 1920 and 2011. The data are shown in Chart 1 (note the large difference in the scales used on the y-axis). It is evident that there are several structural breaks in these series: long periods of tranquillity are occasionally interrupted by bouts of clustered defaults.

Chart 1: Annual issuer-weighted corporate default rates by letter rating, 1920-2011



⁸ The data on bond default should give a good indication of loan default for different rating categories. By construction, our analysis is insensitive to whether this assumption is true: the default process for the hypothetical loan portfolio used to estimate the parameters of the IRB model is constructed on the basis of real-world bond performance.



Source: Moody's Investors Service (2012).

The exercise assumes that the bank holds a large diversified portfolio of corporate loans. The distribution across rating grades AAA to CCC-C is calibrated to match the 2012 Pillar 3 disclosures of Barclays, HSBC, Lloyds Banking Group and RBS, mapping internal risk assessment to external ratings.⁹ Block bootstrap methods are used: the data are split into overlapping blocks of 10 years to preserve autocorrelation and correlation between asset classes. Blocks are then drawn at random, with replacement, to simulate our histories. We simulate 100 alternative histories of default rates on this portfolio, each 100,000 periods long. We do not model rating migration. This simplification does not affect the results. The capital requirement is then estimated using each of the three approaches described in Section 4.1 – Basel I, the Basel II Standardised Approach and the Basel II IRB approach.¹⁰

Implementing the IRB approach requires an estimate of the perceived probabilities of default of each rating grade at each simulation date. Average observed default rates are used for this purpose.¹¹ The look-back period for computing these averages is initially assumed to be a rolling 5 year window, ie the unconditional default probability for B-rated loans is estimated by the simple average default rate experienced by this rating grade in the preceding 5 years.¹² This is the

⁹ In reality, many of the loans will have been made to companies without an external rating, so this is only an approximation.

¹⁰ The term capital is used as a short-hand to include both expected and unexpected losses in the IRB framework. Expected loss is simply $PD \cdot LGD \cdot Exposure$. The Basel I and II standardised approaches do not make this distinction.

¹¹ In practice, banks use various alternative approaches for estimating average default probabilities, including internal default experience, mapping to external data and statistical default models.

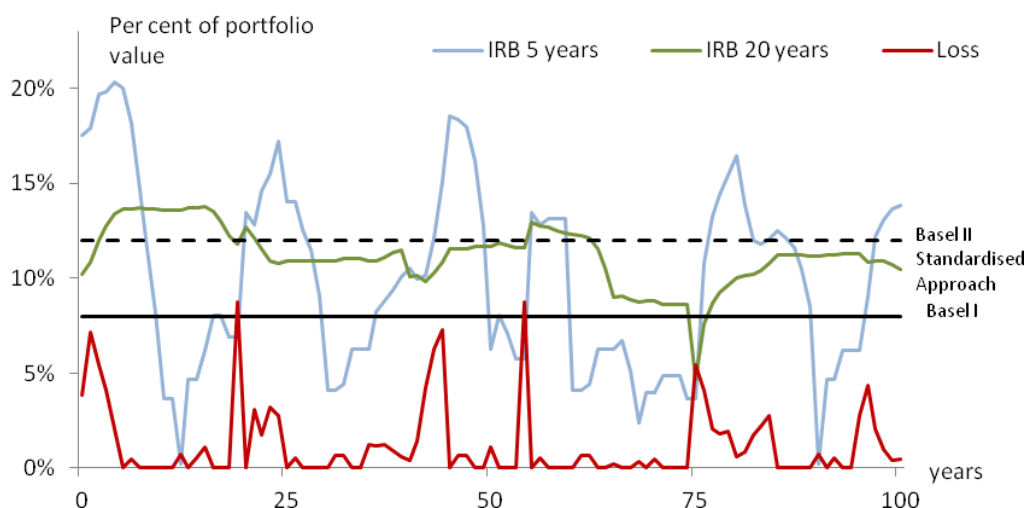
¹² The loan itself is unlikely to have a rating. The rating refers to the rating of the corporate. For simplicity, we refer to these as 'B-rated loans'. In reality a loan from a B-rated company can be relatively more or less risky. In our analysis, we assume that a loan's default probability is the same as that of the issuing entity.

minimum look-back period proposed by the Basel Committee.¹³ The sensitivity of these results to extending this look-back period is then assessed.

4.3 Results

Chart 2 shows an illustrative time series of 100 years of losses (the red line) and capital requirements for a B-rated corporate loan portfolio.

Chart 2: Illustrative losses and capital requirements for B-rated loan portfolio



Source: Authors' calculations.

For this stylised simulation of the IRB approach, required capital is 'high' when there is recent memory of high default rates. But it falls sharply when tranquillity resumes, memories fade, and high default rates drop out of the look-back period used to estimate probabilities of default. If sufficient time passes before defaults unexpectedly pick-up again, banks run the risk of having insufficient capital to absorb losses. When required capital is calculated using a look-back period of 5 years (blue line), there are five violations in this simulation. The fixed risk weights of the Basel I

¹³ See BCBS (2006, paragraph 463). If the available observation period is longer and the data are considered relevant, then the longer period must be used. Firms under the advanced approach are also permitted to estimate the loss given default and exposure at default parameters. The look-back period for these must be no less than seven years, and should ideally cover an economic cycle.

and the Basel II Standardised Approach are robust to the problem of fading memories – though note the Basel I approach does not deliver sufficient capital to prevent violations in periods 19 and 54.

The effectiveness of the IRB approach is therefore likely to depend upon two factors. First, the stability of the process governing defaults: model-based approaches such as IRB are likely to perform well if defaults are a regular occurrence but less well if the defaults are characterised by ‘fat tails’, where long periods of tranquillity are interrupted intermittently and unpredictably by an episode of numerous defaults.

Second, the amount of historical data available for estimating probabilities of default: all else equal, the longer the look-back period, the less sensitive capital requirements will be to recent data, and the smaller the chance of being caught out. To illustrate this point, Chart 2 also plots the IRB capital requirement using a look-back period of 20 years (green line). It is considerably more stable than the 5-year case and, crucially, always delivers significantly more capital in the periods where the 5-year IRB model fails – though it still has one violation in period 75. However, it is not clear *a priori* whether lengthening the look-back period is always desirable. If it is too long, the model may be slow to adapt to structural breaks, for instance.

To explore these ideas more formally, the performance of the alternative approaches to calculating capital requirements is assessed across a large number of simulations. A set of five criteria are proposed to measure performance: the violation rate; cumulative losses; average excess capital; capital variability; and violation clustering. These criteria, described in more detail in Table 2, have been selected to capture the various trade-offs regulators face in practice. All else equal, regulators would prefer a rule that delivered the least number of violations with the least capital, and if there are adjustment costs, with as little variability in the requirement as possible. Where violations do occur, these should be as small as possible and should not be clustered.

Table 2: Criteria for assessing performance

Criterion	Description
Violation rate	The violation rate measures how often losses exceed capital in percentage terms. The IRB formula is calibrated such that this only happens once in a thousand years, so we would expect a violations rate of 0.1%. It is not clear to what level the Basel I and Basel II standardised approaches were calibrated.
Cumulative excess loss	Regulators also care about the magnitude by which a loss exceeds the capital, ie the excess loss. Small violation are less socially costly than large ones. The cumulative excess loss indicator is defined as the sum of losses on average for each 1,000 years of simulation time.
Excess capital	If capital is expensive, the regulatory framework may also seek to avoid requiring banks to have capital in excess of possible realised losses. This indicator is defined as the average of capital over and above losses across the simulation.
Capital variability	If adjusting capital ratios is costly, then, all else equal, regulators will prefer a framework that delivers smooth capital requirements over volatile ones. The coefficient of variation of capital levels (ie the ratio of its standard deviation to its mean) is used to measure capital variability.
Clustering	An underlying assumption of the IRB model is that violations should not be correlated over time. That is, a violation today should not imply that a violation tomorrow is more likely. We test this by comparing the conditional probability of a violation given a violation yesterday to the unconditional probability of a violation (ie the violation rate).

Table 3 reports how well each of the three approaches to calculating capital requirements fares along these five dimensions. To focus on the underlying performance of the IRB model, the 0.03 per cent ‘floor’ for the unconditional default probability is removed in this initial exercise.¹⁴ Higher values indicate poorer performance.

¹⁴ The IRB formula requires a nonzero PD as input. For this reason, a 0.001% floor is retained.

Table 3: Results by asset class

Criterion	Asset class																				
	AA			A			BBB			BB			B			CCC-C			Portfolio		
	IRB	B1	SA	IRB	B1	SA	IRB	B1	SA	IRB	B1	SA	IRB	B1	SA	IRB	B1	SA	IRB	B1	SA
Violation rate	0.024	0	0	0.004	0	0	0.017	0	0	0.031	0	0	0.019	0.012	0	0.064	0.30	0.187	0.01	0.00	0.00
Cumulative excess loss	0.017	0	0	0.010	0	0	0.021	0	0	0.231	0	0	0.362	0.09	0	5.239	24.07	14.33	0.08	0.00	0.00
Excess capital	0.012	0.080	0.016	0.014	0.080	0.040	0.030	0.079	0.079	0.066	0.075	0.075	0.101	0.06	0.104	0.166	0.06	0.088	0.04	0.08	0.06
Capital variability	1.308	n/a	n/a	1.362	n/a	n/a	0.831	n/a	n/a	0.490	n/a	n/a	0.402	n/a	n/a	0.441	n/a	n/a	0.52	n/a	n/a
Clustering	0.000	n/a	n/a	0.000	n/a	n/a	0.000	n/a	n/a	0.026	n/a	n/a	0.241	0.11	n/a	0.882	2.11	1.993	0.00	0.00	0.00
<i>Memo:</i> Excess Kurtosis	9.8			17.2			4.1			15.6			2.7			5.7			2.75		

Notes: B1 stands for the Basel 1 capital requirement; SA stands for the Basel II Standardised Approach; and IRB stands for the Basel II (Foundation) Internal Ratings Based approach. The IRB capital requirements have been calculated using a 5 year look-back period for estimating real time unconditional probabilities of default, with the 3bp floor replaced by a 0.1bp floor. The portfolio is calibrated to be representative of a typical major UK bank's portfolio at the start of the simulation. For each criterion, higher values indicate poorer performance.

There are several noteworthy results. First, at the overall portfolio level, the Basel I and Basel II Standardised approaches both deliver *zero* violations. By contrast, a bank using this stylised simulation of the IRB approach would experience ten times the number of violations this framework is designed to deliver. The required capital for a bank using the simulated IRB approach (on average over the simulation) is, however, only two-thirds that of a bank using the simpler approaches.

Second, while it is performance at the portfolio level that matters, it is nevertheless instructive to explore how the approaches fare for different asset classes with different characteristics, not least because some banks' portfolios will tend to be concentrated in particular segments of the market. For loans rated AA-B, which typically constitute the vast majority of banks' corporate loan portfolios, the Basel I and Basel II Standardised approaches significantly out-perform the simulation of the IRB approach in terms of delivering lower violation rates and cumulative excess losses.¹⁵

The issues are most stark for B-rated loans. Violation rates under the simulated IRB approach are found to be around 50% higher than under Basel I, and the magnitude of losses when they do occur are four times higher. This reflects the phenomenon, illustrated in Chart 2, of model-based capital requirements acting procyclically: they are inappropriately eroded following a temporary period of calm, but then increased sharply following a period of stress. But note that Basel I also achieves this superior performance using only two-thirds as much capital as IRB on average across the simulation. This can also be seen in Chart 2, where IRB overreacts to crises and requires banks to hold considerable excess capital.

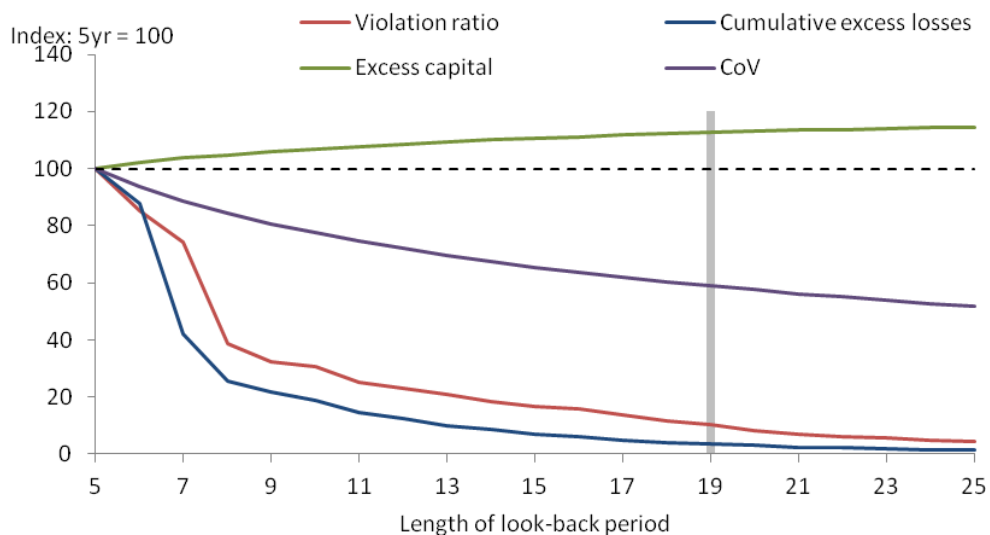
The ranking reverses when we consider the lowest CCC-C-rated loans. For this asset class, the IRB approach delivers significantly fewer excess violations than the other approaches. And the violations, when they do occur, are much larger for the Basel I and II Standardised Approaches. The IRB approach achieves this better performance at the cost of requiring banks to have 2-3 times more capital than Basel I and Standardised approaches.

¹⁵ Based on 2012 Pillar 3 disclosures, AA-B rated corporate loans constituted roughly 98% of major UK banks' corporate loan books assuming a simple mapping from internal ratings to external ratings.

An important driver of these results is the fatness of the tail of the default distribution (ie its kurtosis), which in our model can be thought of as a source of uncertainty. Chart 2 suggested how model-based approaches could be less robust than simple rules in environments with fat-tailed distributions. We investigate this further in two ways. First, we re-scale the Standardised Approach risk weights for each rating grade to deliver the same *average* level of capital across the simulation as the simulated IRB approach. We then re-run the experiments above. We find that the re-scaled Standardised Approach continues to outperform the simulation of the model-based IRB approach. Second, we explore the relative performance of model-based versus simple risk weights within the context of an artificial stochastic process with constant mean and variance, but where kurtosis can be varied. We find that the performance of model-based approaches tends to deteriorate as kurtosis increases.

There are two ways in which this weakness of the IRB approach can be mitigated. The first is to lengthen the look-back period used in the estimation. Chart 3 reports the impact of doing so (leaving out the clustering score as this is zero for all periods). Unless very high weight is placed on the excess capital criterion (the green line), performance is found to improve monotonically as the look-back period increases. For an average corporate loan portfolio, the simulations suggest that roughly 20 years of data are required to achieve the 0.1% violation ratio to which the IRB model is calibrated to (as indicated by the grey bar). As noted above, however, lengthening the look-back period may be unhelpful if the economy is subject to structural breaks.

Chart 3: Performance sensitivity to lengthening the look-back period



Source: Authors' calculations

A second mitigant that enhances the robustness of IRB is the use of simple ‘floors’, which prevent capital requirements falling to excessively low levels after periods of calm. The results reported above abstract from the 0.03% floor for the unconditional default probability that must be used in practice when implementing IRB. When this is introduced, it materially improves the performance of the simulated IRB approach for loans rated BB and above. We explore the benefits of introducing additional floors to these rating categories by using the simple rule of “half the historical probability of default”, which also happens to be 0.03% for AA-rated loans. Doing so considerably improves the performance of the simulated IRB approach: there are no violations from AA-B rated loans, with generally less capital needed than in the Standardised Approaches. For CCC-C rated loans, the floor reduces the number of violations and their magnitude by about a third. This suggests that there may be benefits to a hybrid approach of using models for risk environments, but a simple heuristic (floors) to handle extreme events.¹⁶

Some tentative conclusions can be drawn from this simple exercise. Overall, the results point to specific environments where the use of complex models may be appropriate, such as when they can be based on enough information and when data generating processes are sufficiently stable. But whether such circumstances exist for many real-world bank portfolios is unclear. Our analysis suggests that simpler rule-of-thumb approaches towards risk weighting, such as those provided by the Basel II Standardised Approach for rated portfolios or the even simpler Basel I approach, appear to have better outcomes, though it should be emphasised that they do not address the distinct issue of the serious overall under-capitalisation that emerged under Basel I. Where models are used, they could be made more robust by ensuring they exploit long back runs of data or adding simple rules of thumb such as the capital floor in the IRB formula, which ensures that capital requirements do not fall below a particular threshold. Our results abstract from banks’ responses to the regulatory environment they operate in. An important avenue for further research would be to incorporate ‘risk shifting’ or risk weight ‘optimisation’ effects into the analysis presented above.

¹⁶ Some European banking regulators have recently proposed introducing floors to counteract the secular fall in average risk weights on banks’ mortgage portfolios – mortgage risk weights have fallen as low as 5% for some banks. The Swedish and Norwegian authorities, for instance, have proposed mortgage risk weight floors of 15% and 35% respectively (see Finansinspektionen, 2013 and Norges Bank, 2012).

5 Simplicity versus complexity in the prediction of bank failure

The previous section used a simulation environment to explore the conditions under which simple approaches for calculating capital requirements are likely to succeed or fail relative to complex, model-based ones. This section considers the complementary exercise of whether, empirically, simple indicators and approaches were superior or inferior to more complex ones in predicting failure across a cross-country sample of large banks during the global financial crisis. Two distinct, but related, exercises are conducted.¹⁷ First, individual indicators are analysed in isolation to consider how well simple, $1/N$ -type metrics performed in signalling subsequent bank failure compared to more complex metrics that attempt to weight assets and/or liabilities differently according to their riskiness. Preliminary results from this analysis were presented in Haldane and Madouros (2012); here, a larger set of indicators is considered, including those focussing on the liquidity position of banks.¹⁸

While individual metrics may be useful in their own right, a simple approach does not necessarily equate to a singular focus on one variable, as the discussion of fast-and-frugal trees (FFTs) for medical decision making in Box 1 illustrates (see also Bailey, 2012). So the second exercise considers different approaches for combining the information from individual indicators in trying to predict bank failure. We develop a simple decision tree for assessing bank vulnerability that exploits information from only a small number of indicators via a sequence of binary, threshold rules. The performance of this tree is then compared with commonly used regression-based approaches that attempt to weight all the information optimally (BCBS, 2000; Ratnovski and Huang, 2009; Bologna, 2011; Vazquez and Federico, 2012).

There are several reasons why FFTs might usefully supplement regression-based approaches in improving understanding of bank failure and in communicating risks to relevant stakeholders. First, as the bias-variance trade-off discussed in section 3 illustrates, predictions using less information can sometimes be more robust. Second, since FFTs only

¹⁷ High-level results are presented here; a wider-ranging analysis with further detail is contained in Kapadia, Katsikopoulos and Kothiyal (2014).

¹⁸ Related exercises examining the differential characteristics of global banks which failed during the crisis and those which survived have also been undertaken by IMF (2009, 2011), BCBS (2010) and Arjani and Paulin (2013). But these studies use smaller samples and do not assess explanatory variables in terms of their simplicity; BCBS (2010) is also narrowly focused on capital-based metrics.

assign binary, threshold rules to exploit the information in indicators, they are less sensitive to outliers. Third, FFTs have the advantage that they are able to handle missing data more easily than regression-based approaches and, because they do not weight together different sources of data, they are also more robust to concerns over the reliability or validity of a particular data source. Finally, and perhaps most importantly, while their construction follows an algorithm, the final trees themselves are highly transparent as they provide a clear, simple mapping to the generation of ‘red’ or ‘green’ flags. While such flags would always need to be supplemented by judgement, a tree representation is easier to understand and communicate than the outputs of a regression, whose estimated coefficients give marginal effects.

5.1 *Definitions and data*

The dataset includes almost all global banks which had more than \$100 billion in assets at end-2006.¹⁹ These 116 banks across 25 countries are divided into those that ‘survived’ the crisis (74 banks) and those that ‘failed’ (42 banks), between 2007 and the end of 2009. The definition and classification of failure by and large follows Laeven and Valencia (2010), although this has been supplemented by additional judgement in a few instances.²⁰ A range of balance sheet indicators are collected at consolidated (group) level for each of these banks at end-2005 and end-2006, yielding 11 conceptually distinct indicators, though only a subset

¹⁹ Our list is limited to banks, broker-dealers and building societies / mutuals. We have excluded all federal institutions, diversified financials, speciality lenders and development banks. The institutions are extracted from *The Banker* list of top global banks, supplemented by Capital IQ and SNL databases to extract broker-dealers and building societies / mutuals. Due to data coverage and quality issues, we exclude National Agricultural Cooperative Federation (Korea), Daiwa Securities (Japan) and all Chinese banks from the sample.

²⁰ Because very few banks technically defaulted during the crisis, but many would have without significant government intervention, the definition of failure is necessarily somewhat judgemental. Beyond clear-cut cases of default or nationalisation, Laeven and Valencia (2010) define banks to have failed if at least three of the following six conditions were present: (i) extensive liquidity support (5 percent of deposits and liabilities to non-residents); (ii) bank restructuring costs (at least 3 percent of GDP); (iii) partial bank nationalisation (eg government recapitalisation); (iv) significant guarantees put in place; (v) significant asset purchases (at least 5 percent of GDP); (vi) deposit freezes and bank holidays. In line with IMF (2011), we depart from Laeven and Valencia (2010)’s classifications of Woori Bank, Swedbank and UniCredit as failures, given that the first two benefited from the same market-wide support schemes as other Korean / Swedish banks not defined as failing and UniCredit did not take government assistance despite considering it. We also depart from Laeven and Valencia (2010)’s classifications of Banca Intesa as failing, following the same logic used to classify UniCredit, and Nordea as surviving, given that its 2009 rights issue was partly funded by the Swedish Government and was classed as State Aid under EU law. For the ten institutions in our sample not classified by Laeven and Valencia (2010), we appeal to the classifications of BCBS (2010), IMF (2011) and other publicly available sources, and attempt to follow the Laeven and Valencia (2010) method as closely as possible. Further details are available on request from the authors.

of these are available in the data for some banks.²¹ Table 4 summarises the definition of each indicator; further details are available on request from the authors.

The first metric is the growth rate in total assets (1) between 2005 and 2006, adjusted for significant mergers. When banks are growing very rapidly, they may be prone to move into riskier activities or over-extend themselves in particular markets to achieve such growth. While simple, this metric does not take account of the capital or liquidity resources that banks have available to support their lending, nor does it take into account the overall pace at which the economy is growing.

To gauge capital adequacy, or the ability of banks to withstand losses on their assets without becoming insolvent, data on risk-based Tier 1 capital ratios (2) and leverage ratios (3) are collected. Both have the same numerator, reflecting the amount of equity capital that banks have available to absorb losses. But capital ratios are computed using a measure of risk-weighted assets in the denominator (as discussed in section 4), where less weight is assigned to those assets that are deemed to be less risky. On the other hand, leverage ratios assign all assets the same weight, akin to a $1/N$ rule, and are thus simpler.^{22,23} It should be noted that for the end-2005 and end-2006 data used in this paper, capital ratios were reported on a Basel I basis, with only a few categories of risk, rather than under the Basel II (and Basel III) Standardised or IRB approaches discussed in section 4. Although Basel II aimed to fix some of the flaws of Basel I, it did not address the fundamental under-capitalisation of the regime (a key focus of Basel III), and one conclusion of section 4 is that the greater complexity of risk weighting embedded in Basel II may sometimes lead to worse performance in theory. It is not possible to assess whether this is the case in practice using the type of exercise conducted here.

²¹ These data are primarily collected from SNL and Liquidatum, supplemented by Capital IQ and published accounts in a few instances and Bloomberg for the market-based data.

²² Consistent with regulatory definitions, this paper defines leverage ratios by dividing the relevant measures of capital by assets (eg a leverage ratio of 4%) rather than the reverse (eg a leverage ratio of 25 times).

²³ Due to different accounting standards primarily linked to the treatment of derivatives which permits netting, US and some other banks have a different amount of total assets recorded on their balance sheets from banks elsewhere. Where the data permit, we check the robustness of our results to adjusting such banks' derivative positions in a simple way that should make their total assets figures more comparable – specifically, we add on the difference between the gross value of derivatives and the reported net value. We find that the core results reported throughout section 5 are not materially altered when affected metrics (eg leverage ratios) are calculated using total assets figures which have been adjusted in this way for affected banks.

Indicator	Definition
1 Total asset growth (%) ²⁴	$\frac{\text{Total Assets 2006} - \text{Total Assets 2005}}{\text{Total Assets 2005}} * 100$
2 Basel I risk-based capital ratio (balance sheet) (%)	$\frac{\text{Tier 1 capital}}{\text{Risk weighted assets}}$
3 Leverage ratio (balance sheet) (%)	$\frac{\text{Tier 1 capital}}{\text{Total assets}}$
4 Market-based capital ratio (%)	$\frac{\text{Market capitalisation}}{\text{Risk weighted assets}}$
5 Market-based leverage ratio (%)	$\frac{\text{Market capitalisation}}{\text{Total assets}}$
6 Wholesale funding ratio	$\frac{\text{Wholesale funding level (see 7)}}{\text{Total assets}}$
7 Wholesale funding level ²⁵	Bank deposits + senior paper + collateralised financing (via repo) + wholesale deposits + securitised debt
8 Core funding ratio ²⁶	$\frac{\text{Retail deposits} + \text{long - term wholesale funding} > 1 \text{ yr}}{\text{Total assets}}$
9 Loan to deposit ratio ²⁷	$\frac{\text{Retail loans}}{\text{Retail deposits}}$
10 Net stable funding ratio ²⁸	$\frac{\text{Available stable funding (ASF)}}{\text{Required stable funding (RSF)}}$
11 Liquid asset ratio	$\frac{\text{Cash and balances with central banks} + \text{government bonds}}{\text{Total assets}}$

Table 4: Definitions of variables used

²⁴ Adjusted for significant mergers.

²⁵ This is our preferred measure. For some banks, it is not possible to distinguish between retail deposits and deposits placed by non-bank financial corporations or obtain clean definitions of some of the other components. In these instances, we use close proxies as appropriate.

²⁶ The weighting scheme used to classify different liabilities to determine the core funding ratio on the basis of Liquidatum data is available on request from the authors.

²⁷ This is our preferred measure. For some banks, it is not possible to distinguish between retail deposits and deposits placed by non-bank financial corporations – in these instances, we proxy the loan-to-deposit ratio by (Customer loans / Customer deposits).

²⁸ The weighting scheme used to classify different assets and liabilities to determine the NSFR on the basis of Liquidatum data is available on request from the authors.

Capital and leverage ratios both rely on a regulatory definition of equity capital, which is somewhat opaque. By contrast, market-based capital (4) and leverage (5) ratios use the market capitalisation of banks, as determined in stock markets, while retaining the same denominator as the balance sheet metrics. Arguably, these metrics could be considered as more transparent and thus simpler; on the other hand, the volatility and potential for significant mispricing in financial markets may diminish their usefulness.

We also collect a range of metrics linked to the amount of liquidity risk that a bank faces. To varying degrees, all banks undertake maturity transformation, borrowing at short-term maturities and using the proceeds to finance longer-term loans. But this makes them susceptible to bank runs, whereby large numbers of depositors may simultaneously demand their money back but banks are unable to meet those claims since their assets cannot be liquidated immediately. Several metrics of varying complexity can be used to assess a bank's vulnerability to liquidity risk by capturing the stability of its depositor base, the ease with which its assets may be liquidated, or comparing the two.

Unstable deposits are often those provided 'wholesale' by other financial institutions or capital markets rather than retail deposits. The fraction of wholesale funding in total liabilities (6) therefore provides one very simple measure of liquidity risk. The absolute level of wholesale funding (7) may also be informative in providing a simple gauge on the volume of liabilities that may be particularly flighty, though as a levels measure it is obviously correlated with total assets, so it needs to be interpreted with caution. In the same way that risk weights seek to adjust for the riskiness of different assets, it may be useful to recognise that longer-term wholesale deposits may be more stable than short-term wholesale deposits. The core funding ratio (8) attempts to allow for this by counting long-term wholesale deposits of greater than one-year maturity alongside retail deposits and capital as core funding.

The loan-to-deposit ratio (9) and net stable funding ratio (NSFR) (10) both seek to compare banks' 'stable' funding with their relatively illiquid assets. The former simply views loans as illiquid and compares these to retail deposits so that a high ratio is indicative of a large amount of illiquid assets being financed by wholesale funding and thus high liquidity risk. The latter is a more complex metric in the same vein which weights different asset and liability classes according to their liquidity and stability, with a higher ratio indicative of *lower* liquidity risk. Finally, the liquid asset ratio (11) focusses on the stocks of liquid assets

that banks have readily available to sell or pledge if needed during a period of funding stress.²⁹

In the interests of brevity, the focus in what follows is on the end-2006 observations of the indicators. The first three columns of Table 5 present some key descriptive statistics: the median across all banks for which data points are available for the relevant indicator, and within the sets of failed and surviving banks. Chart 5 presents box plots for a selection of these indicators, separately for banks that failed and for banks that survived. In each case, a box plot provides information about the median value (the horizontal bold line), the 25%- and 75%-percentiles of values and outliers.³⁰ From this preliminary analysis, it is evident that the leverage ratios, measures of wholesale funding and the loan-to-deposit ratio appear to be reasonably good discriminators of subsequent bank failure prior to the global financial crisis, whereas Basel I risk-based capital ratios seem to perform less well.

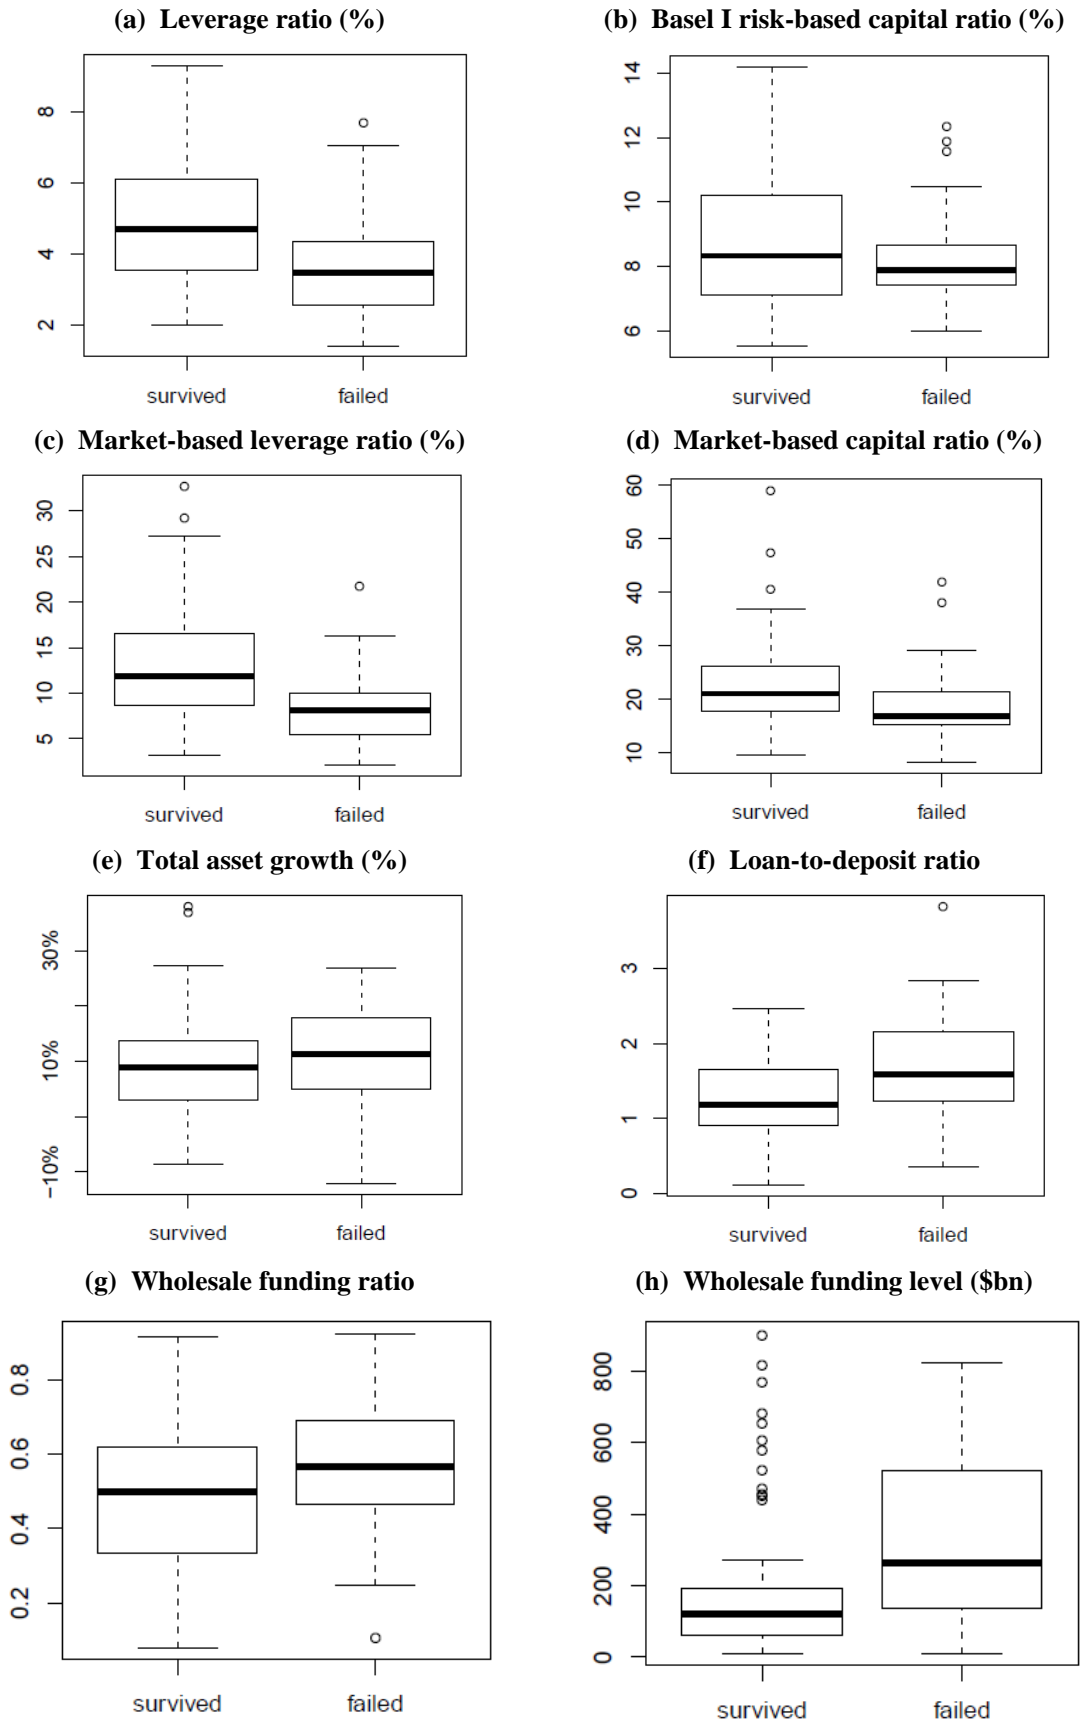
5.2 *Simple or complex indicators: which perform better?*

For our first exercise, we assess more formally the usefulness of each individual indicator in helping to predict subsequent bank failure. To do this, we identify the best cut-off threshold that splits observations of each indicator into zones which give a signal of failure on one side of the threshold and survival on the other. Specifically, a threshold is found for each indicator between its minimum and maximum value in the sample which minimizes the loss function $0.5 \times [\text{Pr}(\text{false alarm}) - \text{Pr}(\text{hit})]$, where $\text{Pr}(\text{hit})$, or the ‘hit rate’, captures the number of banks that are correctly signalled as subsequently failing given the cut-off threshold, relative to the total number of banks that actually failed, and $\text{Pr}(\text{false alarm})$, or the ‘false alarm rate’, captures the number of banks that are incorrectly picked out as subsequently failing given the cut-off threshold, relative to the total number of banks that actually survived.

²⁹ The NSFR is defined under Basel III – see BCBS (2014). We attempt to proxy this definition of the NSFR as closely as possible using data from Liquidatum – further details are available on request from the authors. Basel III also defines a liquidity coverage ratio (LCR) which attempts to compare a bank’s liquid asset buffers with the scale of short-term (less than 30-day maturity) deposits, both weighted in ways designed to capture the differential likelihood that assets may be easily liquidated and deposits withdrawn during periods of stress (BCBS, 2013). But given the complexities of this metric, it is difficult to construct historic LCRs, especially on the basis of publicly available data. The liquid asset ratio measure is also imperfect as some government bonds which count towards it may be encumbered and therefore not available for meeting immediate liquidity needs.

³⁰ A very small number of extreme outliers are excluded to make the core information in the box plots more visible.

Chart 5: Box plots for selected indicators, for surviving and failing banks



This loss function reflects an equal weighting of ‘false alarms’ and ‘hits’. The lower the loss, the better the indicator is – a perfect signal would spot every failure ($\text{Pr}(\text{hit}) = 1$) and yield no false alarms, thus giving a loss of -0.5. As an example, consider the balance-sheet leverage ratio (LR). The minimum and maximum LR across banks in the sample are 1.4% and 9.3% respectively. For any x between these, we have a rule which signals the bank as failing if $\text{LR} < x$ and surviving otherwise. For $x = 4.15\%$, the loss of this rule is minimised at -0.20; hence the cut-off threshold for LR is 4.15%.

We first conduct this exercise by exploiting all of the data available for each individual indicator. This implies that the sample of banks varies slightly across the different indicators considered depending on the extent of data coverage; we briefly discuss our results under a fully consistent sample below. The final three columns of Table 5 give the cut-off thresholds for each metric, indicate whether a value higher or lower than the threshold is the signal of failure, and provide the value of the (minimised) loss function under that threshold. On the whole, the simpler measures tend to perform better. For example, the three best-performing discriminators, the two leverage ratios and the level of wholesale funding, are three of the simplest metrics considered.

In terms of capital adequacy, the balance-sheet leverage ratio performs considerably better than the Basel I risk-based capital ratio, consistent with the findings of IMF (2009). Ignoring the Basel I risk weightings that were supposed to improve measurement increases predictive power in relation to the failure of a typical large bank in this crisis – less is more. It should, however, be noted that the leverage ratio performs relatively less well for the subset of US banks in the sample, which unlike most other banks, were subject to a regulatory restriction on their leverage ratios prior to the global financial crisis. This suggests that the performance of indicators in predicting bank failure may be dependent on the presence or absence of regulatory rules in relation to those metrics.

It is also clear that the market-based capital ratio dominates the balance sheet measure. This could be due to previous deficiencies in the regulatory measurement of the book value of equity, which have partly been resolved by Basel III. But the market pricing of equity may also reflect underlying bank vulnerability more accurately in some circumstances, perhaps because market participants also consider simpler metrics, such as the leverage ratio, or alternative metrics, including in relation to liquidity, when reaching their judgements.

Table 5: Summary statistics and ranking of individual indicators³¹

Indicators	Median	Median of Failed Banks	Median of Surviving Banks	Cut-off Threshold	Failure Signal Location	Minimised Loss
Leverage ratio (balance sheet) (%)	4.16	3.46	4.69	4.15	Lower	-0.20
Market-based leverage ratio (%)	10.06	8.07	11.93	10.14	Lower	-0.20
Wholesale funding level (\$bn)	158	264	122	177	Higher	-0.20
Loan-to-deposit ratio	1.30	1.69	1.19	1.42	Higher	-0.17
Market-based capital ratio (%)	19.74	16.82	21.12	16.82	Lower	-0.15
Wholesale funding ratio	0.51	0.57	0.50	0.34	Higher	-0.11
Basel I risk-based capital ratio (balance sheet) (%)	8.22	7.89	8.34	8.72	Lower	-0.11
Total asset growth (%)	9.6	11.4	8.8	11.2	Higher	-0.10
Net stable funding ratio (NSFR)	0.92	0.91	0.92	1.02	Lower	-0.07
Core funding ratio	0.47	0.44	0.48	0.47	Lower	-0.07
Liquid asset ratio	0.06	0.04	0.06	0.04	Lower	-0.07

Of the liquidity metrics, the structural funding measures that focus on the longer-term liquidity position of banks, especially on the liability side, tend to have the strongest predictive power. For example, the level of wholesale funding and the loan-to-deposit ratio, which both consider the entire liability side of banks' balance sheets, perform very well,

³¹ For each individual indicator, results are reported across all banks for which data points are available for that indicator, so the sample of banks varies slightly across the different indicators considered.

broadly in line with the findings of Ratnovski and Huang (2009), Vazquez and Federico (2012) and Arjani and Paulin (2013) for global banks and Bologna (2011) for US banks. By contrast, the liquid asset ratio, which excludes the liability side altogether, discriminates less well. This may reflect both the difficulty *ex ante* of identifying the relative liquidity of different assets and the fact that liquid assets are typically likely to be insufficient if a bank suffers a sustained outflow of wholesale funding which cumulates to a large fraction of its liabilities. Within the structural funding metrics, it is also striking that the simple wholesale funding and loan-to-deposit measures tend to perform better in relation to this crisis than the more complex metrics such as the NSFR or measures of core funding that attempt to distinguish the maturity and stability of funding and/or assets.

These broad findings continue to apply when we restrict the sample of banks to those for which we have data points for all eleven indicators. The only notable difference is that the minimised loss for some of the best-performing indicators in Table 5 is lowered further i.e. both their absolute and relative signalling power improves when we adopt a consistent sample. But despite these results, it should be emphasised that the indicators which tend to discriminate more poorly on an individual basis in these exercises may still be highly valuable in practice, both because they may still retain important signalling information when taken in conjunction with the better-performing indicators and because future crises may be somewhat different in nature.

5.3 *A fast-and-frugal tree for assessing bank vulnerability*

Although individual indicators can be informative, banks fail for a variety of reasons, suggesting that it is important to combine information from different indicators when trying to assess their vulnerability. While economists typically use regressions for this purpose, we instead start by proposing a FFT similar to those used in medical decision making (see Box 1).

FFTs are characterised by taking a small number of individual indicators, or ‘cues’, with associated thresholds (such as those from the previous subsection), ordering them so that information from one cue is used before moving onto the next, and forcing a classification after each cue via an ‘exit’ from the tree on one of the two sides of the threshold (with the tree continuing to the next cue on the other side of the threshold, except for the last cue, for

which there are exits on both sides). In this context, we think of the classification at the exits as being either a red flag if the bank is judged to be vulnerable or a green flag if the bank is not judged to be vulnerable.

Although FFTs can be constructed statistically, as will be discussed below, we initially construct an intuitive tree for assessing bank vulnerability, shown in Figure 3. Given the need to be parsimonious in the number of indicators used, this tree was constructed by restricting attention to four of the five top-performing individual indicators from Table 5: the balance-sheet leverage ratio, wholesale funding level, loan-to-deposit ratio, and market-based capital ratio. The thresholds for each of these variables are also drawn from Table 5. The ordering and exit structure of the tree is, however, based on economic intuition. But it should be noted that the tree is intended to be purely illustrative of the approach rather than a characterisation of the actual assessment of bank vulnerability, for which judgement must always play a key role.³²

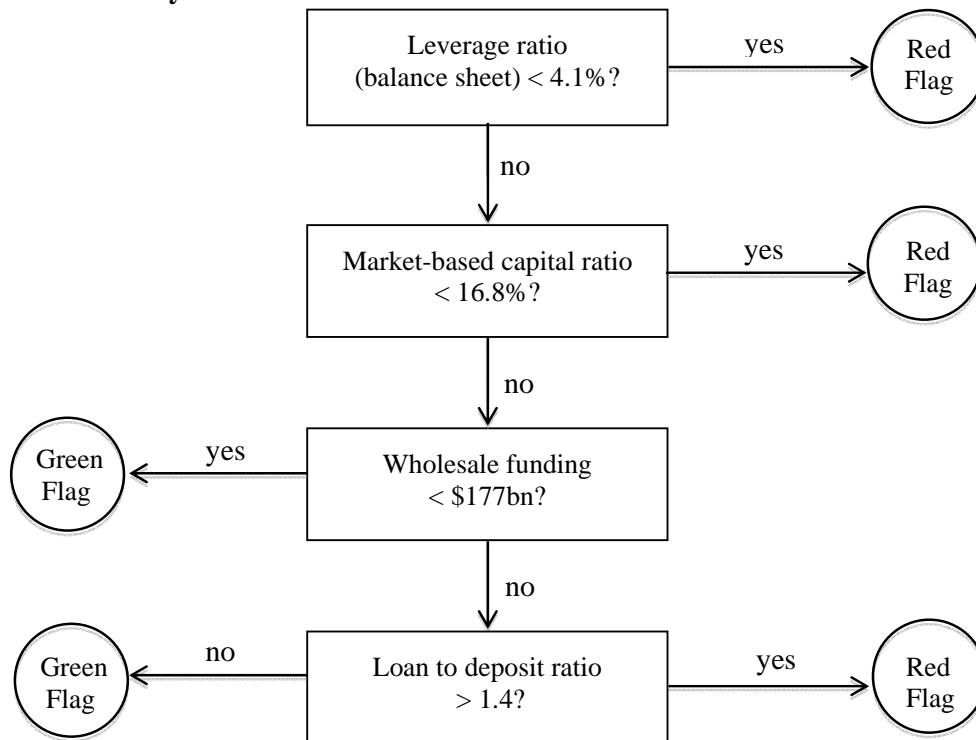
We adopt the balance-sheet leverage ratio as our first cue. While this choice is partly driven by its performance as the (joint) best-discriminating individual indicator, it also reflects the fact that the leverage ratio is a simple measure of the underlying solvency of banks, giving a direct read on its likelihood of eventual failure, even if that failure is precipitated by a liquidity crisis. At the same time, it would not seem sensible to give banks a green flag solely on the basis of having a high leverage ratio; therefore, the exit instead automatically gives a red flag to all banks with a leverage ratio of less than 4.1%.

A significant drawback of leverage ratios is that, in isolation, they do not penalise banks in any way for the riskiness of their assets (Tucker, 2012; Carney, 2012; Bailey, 2012). While we have argued that risks may be difficult to compute accurately, their ordinal ranking may sometimes be relatively straightforward to judge – for example, a mortgage with a loan-to-value ratio of 95% is typically likely to be more risky than a mortgage with a loan-to-value ratio of 25%. In view of this, risk-based capital ratios, which have the capacity to penalise more risky assets, are likely to be an important complement to leverage ratios. So, it makes sense to put the market-based capital ratio as the second cue and use it to assign a bank a red flag if it signals vulnerability, even if the bank has a leverage ratio of above 4.1%. By

³² See Bank of England and Financial Services Authority (2012).

contrast, giving any bank a green flag solely based on these two capital-based metrics would seem dangerous given the central contribution of liquidity weaknesses to many bank failures.³³

Figure 3: An example judgment-based (intuitive) fast-and-frugal tree for assessing bank vulnerability



The third and fourth cues, therefore, turn to liquidity metrics. Since a bank’s wholesale funding level is another of the (joint) best discriminating individual indicators, we use this as the third cue; as noted above, however, this indicator needs to be interpreted with caution given that it is a levels indicator not normalised by balance sheet size, so in the analysis which follows we also report results excluding this indicator. For this cue, we exit to a green flag if a bank’s wholesale funding is below a particular level on the grounds that most liquidity crises arise from wholesale funding (at least initially), so a bank is less likely to be at risk if it has a small volume of wholesale funding; a more conservative regulator might

³³ Despite its strong individual performance, we do not additionally consider the market-based leverage ratio for the FFT given that it shares some common features with both the balance-sheet leverage ratio and market-based capital ratio.

instead choose to exit to a red flag on the other side of the threshold. Finally, wholesale funding is more likely to be a concern if it is used to finance very illiquid assets. As noted above, this can be partially captured by the loan-to-deposit ratio, our final cue. If this is below 1.4, we give the bank a green flag; otherwise we give it a red flag.

To give an idea of how this tree might work in practice, consider the case of UBS, which required significant support from the Swiss authorities during the crisis. As it had a leverage ratio of 1.7% at end-2006, it is automatically given a red flag at the first cue in the tree. This is despite the fact that it had a market-based capital ratio significantly exceeding 16.8% and a loan-to-deposit ratio well below 1.4 – the FFT completely ignores this information. By contrast, a regression would balance UBS’s low leverage ratio with its high market-based capital ratio and low loan-to-deposit ratio and, therefore, perhaps not give such a strong red signal.

On the other hand, the FFT does not successfully identify Wachovia as vulnerable based on pre-crisis data. With a leverage ratio of 5.6% and a market-based capital ratio of 20.4%, it is not given a red flag on the basis of the capital-based metrics. Nor, however, is it given a green flag at the third cue due to its high level of wholesale funding. But its loan-to-deposit ratio of 1.21 gives it a green flag at the final cue. As well as illustrating how the tree may be used, this example highlights how judgement must always remain central to any assessment of bank vulnerability – in Wachovia’s case, its particularly large subprime exposures contributed to its downfall.

More generally, we can calculate the overall in-sample performance of this tree across the entire dataset of banks. Doing this, we find that it correctly calls 82% of the banks that failed during the crisis (a hit rate of 0.82), while, of the total number of banks that survived, 50% of those were incorrectly called as failing (a false alarm rate of 0.50). Excluding the wholesale funding level indicator from the tree increases the hit rate to 0.87 but the false alarm rate goes up to 0.63. Together, these results suggest that the tree has reasonably good in-sample performance, especially in assigning red flags to banks which subsequently failed.³⁴

³⁴ Note that if hits and false alarms are weighted equally in the loss function, the balance-sheet leverage ratio indicator taken individually has a slightly lower loss than both of these trees, primarily because its lower hit rate (0.74) is more than outweighed by a lower false alarm rate (0.35). This would seem to argue for using individual indicators rather than the trees. There are, however, two reasons opposing this argument. First, if

While intuitive approaches may be very useful, it is also interesting to consider how FFTs for assessing bank vulnerability may be constructed via an algorithm (see also Martignon et al, 2008 and Luan, Schooler and Gigerenzer, 2011). To do this, we take the same three / four individual indicators used above but order the cues according to their loss for the function $0.5 \times [\text{Pr}(\text{false alarm}) - \text{Pr}(\text{hit})]$ (as used for the ranking of individual indicators in the previous subsection) rather than selecting their order based on intuition. The other departure from the intuitive tree of Figure 3 is that the algorithm makes a decision about what kind of exit – red flag or green flag – to place at each level of the tree, taking the threshold values as given and fixed at the values which minimise loss for the indicator on an individual basis. The algorithm classifies the exits by exhaustively enumerating all possibilities and choosing a tree structure that delivers the best overall categorization performance according to the loss function $0.5 \times [\text{Pr}(\text{false alarm}) - \text{Pr}(\text{hit})]$.³⁵

The above exercise is then repeated under a range of different loss functions which move beyond an equal weighting of false alarms and hits. Specifically, we adopt the more general loss function, $w \times \text{Pr}(\text{false alarm}) - (1 - w) \times \text{Pr}(\text{hit})$, and apply the algorithm under a wide range of w , including $w = 0.5$ (which is the base case above). In undertaking this exercise, we retain focus on the same three / four indicators but their loss and thus ordering, and associated thresholds, vary with w .

We restrict the sample to only those banks which have data points for all four indicators – while the ability to handle missing data is a strength of FFTs, this facilitates subsequent comparability with a regression-based approach. To avoid over-fitting, we do not estimate the thresholds of the indicators, and associated losses and ordering, or the sequence of exits based on the entire sample, but instead use a training set containing 70% of the banks in the

higher weight were placed on the hit rate than the false alarm rate in the loss function, then the trees would start to perform better. Second, there are strong economic arguments as to why a narrow focus on a single indicator might be undesirable – one may miss obvious risks that might be building in other areas not captured by the indicator in question or it may create significant arbitrage opportunities.

³⁵ Rather than using the performance of individual indicators to fix the overall ordering of the indicators and all of the threshold values from the start, these could be adjusted endogenously as the tree progresses depending on the classification already given by previous indicators further up the tree. In other words, loss functions and thresholds could be recomputed for remaining banks in the sample not already classified by cues earlier in the tree. It should be noted, however, that such computationally intensive approaches run the danger of over-fitting the data (Gigerenzer and Brighton, 2009). A host of such approaches are explored in Kapadia, Katsikopoulos and Kothiyal (2014).

sample, randomly chosen. The performance of the resulting tree is evaluated against the remaining 30% of banks. This process of estimating and evaluating is repeated 1,000 times using different training sets to average out random variation and is an important feature of the approach. The result is a sequence of points, each corresponding to a different w , which give the average hit rate and false alarm rate over all trees constructed with that w . Such sequences of hit and false alarm rates are commonly known as ‘receiver operating characteristic’ (ROC) curves.


Figure 4 presents the ROC curves under this algorithmic approach for constructing FFTs, both including and excluding the wholesale funding level indicator (panels (a) and (b) respectively). The x -axis corresponds to the false alarm rate and the y -axis to the hit rate. Thus, performance increases as points move towards the upper-left hand corner, as this region maximises hits while minimising false alarms. The figures also include the out-of-sample performance of intuitive trees corresponding to Figure 3 but with updated thresholds based on the restricted, consistent sample. These are represented by a single point in each panel. Interestingly, they perform slightly better than the corresponding trees generated by the algorithm.

While it is evident that the FFTs have some predictive power, a key question is how their performance compares to regression-based approaches. To assess this, we use the same three / four variables as inputs into a logistic regression, again selecting random 70% selections of the sample to fit the model and using the remaining 30% for out-of-sample prediction.³⁶ Under this approach, the output for each bank is a probability p of bank failure, which is transformed to a decision by using another parameter t : if $p > t$, then the bank is given a red flag; otherwise it is given a green flag. The trade-off between the hit rate and the false alarm rate may be controlled by varying t and we obtain a ROC curve by choosing values of t which minimise loss for the same values of w used in the construction of the FFT ROC curve.


³⁶ We also cross-check our results using a Probit approach to constructing the regression model, finding similar results. In future work, it would be interesting to consider broader methods including the use of a linear probability model, possibly under a quantile regression approach.

Figure 4: Receiver operating characteristic curves for (i) a logit model, (ii) a statistical fast-and-frugal tree and (iii) a judgement-based (intuitive) fast-and-frugal tree

(a) with four indicators (including wholesale funding level indicator)



(b) with three indicators (excluding wholesale funding level indicator)



Plotting the regression results in Figure 4, we find that the logistic regression performs comparably to the statistical FFT. When the wholesale funding level indicator is included, it outperforms the statistical FFT (though not the intuitive FFT) but when that indicator is excluded, the statistical FFT tends to do better. The difference across the two exercises highlights the importance of considering a range of other tests, including further varying the number of indicators used, for example by pruning the trees and regularisation for the regression, considering different methodologies for constructing FFTs, and using a wider range of data before reaching firm conclusions on the relative predictive performance of each method in this context.

5.4 *Discussion*

These preliminary results highlight the potential usefulness in regulators using simple indicators to complement their assessment of individual bank vulnerability. In particular, simple indicators seemed to perform better than more complex ones in signalling subsequent failure across a cross-country sample of large banks during the global financial crisis. The results also highlight how there can be situations in which a FFT, which ignores some of the information by forcing an exit based on a binary, threshold rule at each stage and does not attempt to weight different indicators together, can perform comparably to more standard regression-based approaches. The FFT is, arguably, also more intuitive and easier to communicate, providing a quick and easy schematic that can be used to help classify banks. For example, while it would always need to be supplemented by judgement, it could be used as a simple, intuitive means of explaining why certain banks may be risky even if the management of those banks argue that they are not taking excessive risks, or as one method for informing which institutions supervisory authorities should focus their efforts on from the very large set for which they typically have responsibility.

As discussed above, however, these findings may partly be a product of the regulatory regime in place during the period under investigation. This both emphasises the risks from regulatory arbitrage and other adverse incentive effects that would arise from focussing on just a single indicator such as the leverage ratio and highlights the possibility that indicators which appeared to signal well in the past may lose some of their predictive power when they become the subject of greater regulatory scrutiny. More generally, when interpreting

indicators that were useful in the past, it is important to recognise that future crises may be somewhat different in nature and judgement must always play a key role in assessing risks.³⁷

Subject to these caveats, the results may have lessons for the design of regulatory standards. Together with the findings from section 4, they highlight the importance of imposing a leverage ratio standard to complement risk-based capital requirements. And the predictive performance of simple structural funding metrics emphasises the importance of reaching international agreement in this area on suitable microprudential standards, possibly with simple metrics complementing the proposed NSFR, at least in terms of assessing liquidity risks.

While an area for future analysis, the extension of these ideas to the macroprudential sphere may have important implications for both the design of instruments and macroprudential risk assessment. In particular, it speaks to the use of simple instruments for macroprudential purposes. And it suggests using simple, high-level indicators to complement more complex metrics and other sources of information for assessing macroprudential risks (Aikman *et al*, 2013; Bank of England, 2014). More generally, simplicity in macroprudential policy may also facilitate transparency, communicability and accountability, thus potentially leading to a greater understanding of the intent of policy actions, which could help reinforce the signalling channel of such policies.³⁸

6 Conclusion

This paper has argued that financial systems are better characterized by uncertainty than by risk because they are subject to so many unpredictable factors. As such, conventional methods for modeling and regulating financial systems may sometimes have drawbacks. Simple approaches can usefully complement more complex ones and in certain circumstances less can indeed be more. This is borne out to a degree by both simulations of capital requirements against potential losses and the empirical evidence on bank failures during the global financial crisis, with potentially important lessons for the design of financial regulation.

³⁷ See Goodhart (1975) and Lucas (1976).

³⁸ See Giese *et al* (2013).

It may be contended that simple heuristics and regulatory rules may be vulnerable to gaming, circumvention and arbitrage. While this may be true, it should be emphasised that a simple approach does not necessarily equate to a singular focus on one variable such as leverage – for example, the FFT in section 5 illustrates how simple combinations of indicators may help to assess bank vulnerability without introducing unnecessary complexity. Moreover, given the private rewards at stake, financial market participants are always likely to seek to game financial regulations, however complex they may be. Such arbitrage may be particularly difficult to identify if the rules are highly complex. By contrast, simpler approaches may facilitate the identification of gaming and thus make it easier to tackle.

Under complex rules, significant resources are also likely to be directed towards attempts at gaming and the regulatory response to check compliance. This race towards ever greater complexity may lead to wasteful, socially unproductive activity. It also creates bad incentives, with a variety of actors profiting from complexity at the expense of the deployment of economic resources for more productive activity. These developments may at least partially have contributed to the seeming decline in the economic efficiency of the financial system in developed countries, with the societal costs of running it growing over the past thirty years, arguably without any clear improvement in its ability to serve its productive functions in particular in relation to the successful allocation of an economy's scarce investment capital (Friedman, 2010).

Simple approaches are also likely to have wider benefits by being easier to understand and communicate to key stakeholders. Greater clarity may contribute to superior decision making. For example, if senior management and investors have a better understanding of the risks that financial institutions face, internal governance and market discipline may both improve. Simple rules are not a panacea, especially in the face of regulatory arbitrage and an ever-changing financial system. But in a world characterised by Knightian uncertainty, tilting the balance away from ever greater complexity and towards simplicity may lead to better outcomes for society.

References

- Aikman, D, Barrett, P, Kapadia, S, King, M, Proudman, J, Taylor, T, de Weymarn, I and Yates, T (2011), Uncertainty in macroeconomic policy-making: art or science? *Philosophical Transactions of the Royal Society A*, vol. 369, no. 1956, pages 4798-4817.
- Aikman, D, Haldane, A and Kapadia, S (2013), Operationalising a macroprudential regime: goals, tools and open issues, *Banco de España Financial Stability Review*, 24, 9-30.
- Anderson, R. and May, R. (1991). *Infectious diseases of humans: dynamics and control*. Oxford and New York: Oxford University Press.
- Arjani, N and Paulin, G (2013), 'Lessons from the financial crisis: bank performance and regulatory reform', *Bank of Canada Discussion Paper*, 2013-4.
- Arrow, K. J. and Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica* 22: 265–290.
- Bailey, A (2012), 'The challenges in assessing capital requirements for banks', speech given at the Bank of America Merrill Lynch conference, London, 6 November.
- Bank of England and Financial Services Authority (2012), 'The PRA's approach to banking supervision', PRA approach documents, October.
- Bank of England (2014), 'The Financial Policy Committee's powers to supplement capital requirements: a Policy Statement', January.
- Benartzi, S. and Thaler, R.(2001), 'Naive diversification strategies in defined contribution saving plans', *American Economic Review*, American Economic Association, 91(1), 79-98.
- Bikhchandani, S, Hirshleifer, D and Welch, I (1992), 'A theory of fads, fashion, custom, and cultural change as informational cascades', *Journal of Political Economy*, 100(5), pages 992-1,026.
- Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy* 81 (3): 637–654.
- BCBS (2000), Credit ratings and complementary sources of credit quality information, *Basel Committee on Banking Supervision Working Papers*, No. 3.
- BCBS (2005), 'An Explanatory Note on the Basel II IRB Risk Weight Functions', July.
- BCBS (2006), 'Basel II: International Convergence of Capital Measurement and Capital Standards: A Revised Framework – Comprehensive Version', June.
- BCBS (2010), 'Calibrating regulatory minimum capital requirements and capital buffers: a top-down approach', October.
- BCBS (2013), 'The liquidity coverage ratio and liquidity risk monitoring tools', January.

- BCBS (2014), 'Basel III: The Net Stable Funding Ratio: Consultative Document', January.
- Bologna, P (2011), 'Is there a role for funding in explaining recent US banks' failures?' *IMF Working Paper*, 11/180.
- Bower, B. (2011), Simple heresy: Rules of thumb challenge complex financial analyses. *Science News*, 179, 26–29.
- Brighton, H. and Gigerenzer, G. (2012). How heuristics handle uncertainty. In P. M. Todd, G. Gigerenzer and the ABC Research Group, *Ecological rationality: Intelligence in The World* (pp. 33-60). New York: Oxford University Press.
- Caballero, R.J. and Pindyck, R.S. (1996), 'Uncertainty, Investment, and Industry Evolution', NBER Working Paper No. w4160.
- Cacciolo, F. Marsili, M., and Vivo, P. (2009), 'Eroding market stability by proliferation of financial instruments', *European Physical Journal B*, 71, 467-479.
- Camerer C. (1999). Behavioral economics: Reunifying psychology and economics. *Proceedings of the National Academy of Sciences of the United States of America*, 96, 10575-7.
- Carney, M (2012), 'Some current issues in financial reform', remarks at the Canadian Club of Montréal, Montréal, 8 November.
- Czerlinski, J., Gigerenzer, G. and Goldstein, D.G. (1999). How good are simple heuristics? In Gigerenzer, G, Todd, P M and the ABC Research Group, *Simple heuristics that make us smart* (pp. 97-118). New York: Oxford University Press.
- Derman, E. and Wilmott, P. (2009), *The Financial Modelers' Manifesto*, Wilmot.com.
- DeMiguel, V., Garlappi, L. and Uppal, R. (2007). Optimal versus naïve diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22, 1915-1953.
- Finansinspektionen (2013), Risk weight floor for Swedish mortgages. See http://www.fi.se/upload/90_English/20_Publications/20_Miscellaneous/2013/riskviktsgolvsvenska-bolan-12-11920-21maj2014-eng.pdf.
- Friedman, B M (2010), 'Is our financial system serving us well?' *Daedalus*, 139(4), pages 9-21.
- Gai, P., Haldane, A.G. and Kapadia, S. (2011), Complexity, concentration and contagion, *Journal of Monetary Economics*, Vol. 58, Issue 5, pages 453-470.
- Gai, P. and Kapadia, S. (2010), Contagion in Financial Networks, *Proceedings of the Royal Society A*, vol. 466, no. 2120, pages 2401-2423.
- Geman, S., Bienenstock, E. and Doursat, E. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1-58.

- Giese, J., Nelson, B., Tanaka, M., Tarashev, N. (2013), How could macroprudential policy affect financial system resilience and credit? Lessons from the literature, *Financial Stability Paper No 21*, Bank of England.
- Gigerenzer, G. and Selten, R. (Eds.) (2001). *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
- Gigerenzer, G. and Brighton, H. (2009). Homo heuristics: Why biased minds make better inferences. *Topics in Cognitive Science, 1*, 107-143.
- Gigerenzer, G., Hertwig, R. and Pachur, T. (Eds.) (2011). *Heuristics: The Foundations of Adaptive Behavior*. Oxford: Oxford University Press.
- Goodhart, C (1975), Problems of monetary management: the U.K. experience, *Papers in Monetary Economics*, Reserve Bank of Australia.
- Green, L.A. and Mehr, D.R. (1997). What alters physicians' decisions to admit to the coronary care unit? *Journal of Family Practice, 45*, 219-226.
- Haldane, A. G. and Madouros, V. (2012), The Dog and the Frisbee, Speech given at the Federal Reserve Bank of Kansas City's 36th Economic Policy Symposium, Jackson Hole, 31 August.
- Hansen, M.H., Hurwitz, W.N. and Madow, W.G. (1953). *Sample Survey Methods and Theory*. New York: John Wiley & Sons, Inc.
- Hansen, L.P. and Sargent, T.J. (2007), *Robustness*, Princeton: Princeton University Press.
- International Monetary Fund (2009): *Global Financial Stability Report*. Chapter 3: Detecting Systemic Risk.
- International Monetary Fund (2011): *Global Financial Stability Report*. Chapter 2: How to Address the Systemic Part of Liquidity Risk.
- Kahneman D. and Tversky A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*, 263-291.
- Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Kapadia, S., Katsikopoulos, K. and Kothiyal, A. (2014), TBA
- Katsikopoulos, K. V. (2011). Psychological heuristics for making inferences: Definition, performance, and the emerging theory and practice, *Decision Analysis, 8*, 10-29.
- Katsikopoulos, K. V., Schooler, L. J., and Hertwig, R. (2010). The robust beauty of mediocre information. *Psychological Review, 117*, 1259-1266.
- Knight, F. H. (1921). *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin.
- Laeven, L. and Valencia, F. (2010), 'Resolution of Banking Crises: The Good, the Bad and the Ugly', *IMF Working Paper*, 10/146.

- Luan, S., Schooler, L. J. and Gigerenzer, G. (2011). A signal-detection analysis of fast-and-frugal trees. *Psychological Review*, 118, 316-338.
- Lucas, R (1976), Econometric Policy Evaluation: A Critique, in Brunner, K.; Meltzer, A., *The Phillips Curve and Labor Markets*, Carnegie-Rochester Conference Series on Public Policy 1, New York: American Elsevier, pp. 19–46,
- Makridakis, S. and Hibon, M. (1979). Accuracy of forecasting: an empirical investigation (with discussion). *Journal of the Royal Statistical Society A* 142, 97-145.
- Makridakis, S. and Hibon, M. (2000) The M3-competition: results, conclusions and implications. *International Journal of Forecasting*, 16, 451-476.
- Mandelbrot, B.B. and Hudson, R.L. (2004). *The (Mis)Behavior of Markets: A Fractal View of Risk, Ruin, and Reward*. New York: Basic Books.
- Markowitz, H. M. (1952), Portfolio Selection, *Journal of Finance*, 7 (1), pp 77-91.
- Martignon, L., Katsikopoulos, K. V. and Woike, J. K. (2008). Categorization with limited resources: A family of simple heuristics. *Journal of Mathematical Psychology*, 52, 352-361.
- Martignon, L. and Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision*, 52, 29-71.
- May, R. and Arinaminpathy, N. (2010) . Systemic risk: the dynamics of model banking systems. *Journal of the Royal Society Interface* 46, 823–838.
- Merton, R. C. (1969), Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case, *Review of Economics and Statistics*, 51 (3), August, pp 247-57.
- Moody’s Investors Service (2012), Corporate default and recovery rates, 1920-2011.
- Mugford, S. T., Mallon, E. B. and Franks, N. R. (2001). The accuracy of Buffon’s needle: A rule of thumb used by ants to estimate area. *Behavioral Ecology*, 12, 655-658.
- Norges Bank (2012). *Financial Stability Report*, November.
- Petrie, M. and Halliday, T. (1994). Experimental and natural changes in the peacock’s (*Pavo cristatus*) train can affect mating success. *Behavioral Ecology and Sociobiology*, 35, 213-217.
- Rabin M. (1998). Psychology and economics. *Journal of Economic Literature*, 36, 11-46.
- Ratnovski, L and Huang, R (2009), Why are Canadian banks more resilient? *IMF Working Paper*, WP/09/152.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99-118.

Subbarao, D (2011), 'India and the global financial crisis: what have we learnt?', speech given at the South Asia Research Centre of the Australian National University, Canberra, 23 June.

Taguchi, G and Clausing, D (1990), 'Robust quality', *Harvard Business Review*, 68(1), 65-75.

Taleb, N. (2007). *The Black Swan: The Impact of the Highly Improbable*. New York: Random House.

Tarullo, D.K. (2008), *Banking on Basel: The Future of International Financial Regulation*. Washington D.C: Peterson Institute for International Economics.

Tucker, P (2012), Competition, the pressure for returns, and stability. Speech given at the British Bankers' Association Annual Banking Conference, London, 17 October.

Vazquez, F, and Federico, P. (2012), 'Bank Funding Structures and Risk: Evidence from the Global Financial Crisis'. *IMF Working Paper*, 12/29.

Wübben, M. and Wangenheim, F. v. (2008). Instant customer base analysis: Managerial heuristics often "get it right." *Journal of Marketing*, 72, 82-93.