BANK OF ENGLAND

Speech

Analytical Models of Financial Stability

Speech given by
Alastair Clark, Adviser to the Governor, Bank of England

At the Cass Business School
28 March 2007

Accompanying Slides

1 [Introduction]

2 What I want to talk about this evening is work in progress on two levels. First, the topic – analytic models of financial stability – is an area of active interest in the Bank and elsewhere; but I do not think anyone would claim to have progressed beyond the early stages of developing such models. Indeed, given the complexity of the underlying issues, it may be some time before we have a really satisfactory quantitative framework for analysing financial stability questions. But that does not mean we should not try. Second, this talk is itself work in progress – more so than I hoped when I originally agreed to give it! Leaving the Bank has proved a more time-consuming business than I had anticipated. Nevertheless, I thought it might be interesting to sketch out some of the challenges, as we see them, in trying to understand what makes for a stable financial system, and in thinking about what public policy can contribute to financial stability.

Background

3 Let me start with a few general observations.

4 First, it is now conventional to recognise, but nonetheless true, that when we refer to "financial stability" we are often not entirely clear what we're talking about – or more accurately that different people may be talking about different things. First-hand evidence of this is provided on the financial stability web site run by my erstwhile colleague Phil Davis. Some define it, or rather its obverse "financial instability", in very broad terms to mean defects in the mechanism for allocating savings to investment. To my mind, this goes too wide and potentially includes all sorts of inefficiencies in financial intermediation which do not, in common sense usage, constitute instability. But in the reverse direction, it is of course true that instability can be an important source of inefficiency. Others have defined financial instability in terms of volatility, or "excessive" volatility, in asset prices. From this perspective, the mechanics of the financial system may function perfectly satisfactorily at a micro level; the problem, if there is one, lies in the dynamic characteristics of the system as a whole. Perversely, it may be that the more successful we are in removing frictions in financial markets and in institutional

behaviour, the more likely it is that instability in this sense will occur. Allied to this, a third approach is to draw an analogy with "stability" in the physical sciences – that is to view the financial system as a kind of giant machine, and focus on the structural characteristics which determine whether, when the system is disturbed, it tends to dampen the disturbances and return to its starting point or fly off in some new direction, or collapse. And finally, and perhaps closest to the "popular" view if such exists, financial instability can be defined as arising when there is some substantial discontinuity in the functioning of the financial system – institutional insolvency, market unavailability, or whatever – on a scale which leads to significant costs for the economy as a whole.

5 These various definitions clearly have elements in common and indeed they all have merits in their own terms. My purpose in rehearsing them is not to recommend a "best buy". It is simply to emphasise that in constructing analytic models of financial stability it is important to be clear which definition of stability is in mind.

6 A second general point is that, in contrast to its twin, monetary stability, financial stability is not easily measured. You could argue about inflation rates, exchange rates or monetary aggregates as optimal indicators of monetary conditions – although you know where the Bank's own preference lies. All of them are plausible candidates and, crucially, each of them can be rooted in a model of how monetary conditions relate to the wider economy and of how the transmission mechanisms for monetary policy work. But as yet we have nothing similar for financial stability and as a result no generally accepted metrics for determining whether the financial system has become more or less stable. For those whose job it is to try to increase the resilience of the financial system, this is a considerable handicap.

7 Third, specifically, it is important to be clear in talking about the potential welfare costs of some external shock – external that is to the financial system – whether the costs associated with the shock itself are to be included, or only the additional costs incurred through so-called "amplification factors" – frictions of various kinds – as the effects of the shock percolate through the system. In principle the latter is probably the more sensible measure although in practice the two elements may be difficult to disentangle.

8 A fourth general point is that the objective of financial stability analysis is, in the broadest terms, to try to understand what is likely to happen in the tails of probability distributions. It is unlikely that a financial crisis will develop if everything evolves in the way most economic agents expect – for the simple reason that if they see financial stresses coming they will probably take avoiding action. This may not always be true; coordination problems may mean that adjustment is too little or too late. But serious problems are more likely if the unexpected happens. Behaviour in the tail of distributions is often not well understood and there is much less data with which to try to analyse it.

9 Finally, and a related point, speed is of the essence in the development of financial instability, certainly in the "discontinuity" sense. Even large shocks can be accommodated if there is plenty of time to adjust – we saw this for example in the reaction of the US banking system to the Latin American debt crises of the 'eighties and more recently in the reaction to the Argentine default. The position becomes much more dangerous when shocks come in rapid succession and contagion effects build up, so that decisions have to be made with inadequate – probably increasingly inadequate – information and the normal mechanisms for redistributing or absorbing losses are overwhelmed. In military terminology, the defences become saturated. Unfortunately, it is precisely these timing effects which are amongst the most difficult to model, the more so since they depend importantly on confidence which is itself notoriously difficult to capture analytically.

10 In the face of these challenges, the prudent central banker, and perhaps even the prudent academic, might be inclined to say that it all looks very interesting - and then pass swiftly on. There are, however, some powerful reasons for trying to do better than that. In the UK we have not experienced a serious financial crisis for a long time. But other countries, and not just emerging markets, have: one only needs to look to Scandinavia and Japan for recent examples. And the costs in terms of foregone GDP can be large, certainly as large as those arising from macroeconomic under-performance in a more conventional sense – although they tend to come in big but infrequent dollops and for that reason are easily forgotten or underestimated. Quantifying the losses is difficult but estimates in the Bank suggest that 25% of GDP

would not be unreasonable as a typical order of magnitude. Understanding how public policy might contribute to the avoidance of such costs therefore seems a worthwhile objective.

11 At a more day-to-day level, too, it would be helpful to have, for example, a better way of assessing how high-level regulatory changes are likely to affect the overall robustness of the financial system, partly so as to inform a cost/benefit analysis of such measures and partly as a guide in ranking different policy options. And again, a better model of institutional interactions in the financial sector would help in understanding how the risks associated with individual firms relate to the overall stability of the system. For all these and other reasons, the Bank has devoted significant resources over recent years to the goal of developing models of the financial sector which might improve our analysis of its stability.

Categories of models

12 Dividing up models of financial stability into a set of hard-and-fast categories is probably rather artificial; there are no hard boundaries. Models based on one approach can be expanded or modified to incorporate elements of other approaches; and certainly models based on different approaches can usefully complement one another. For convenience, however, I will say a few words about these models under three headings, which I have labelled: augmented macro-economic models; dynamic stochastic general equilibrium models based on micro optimising behaviour; and network models. [**Slide 1.]**

Augmented macro models

13 Augmented macro models take the outputs of a "conventional" macro model, the kind we are all familiar with, and feed them into a set of relationships which aim to capture the impact on key financial variables. There may then be further links in the chain so that changes in these financial variables can be traced through to impacts on, say, the balance sheets of financial intermediaries. A key feature is that, rather than looking at point estimates of the variables concerned, they are treated as stochastic. Any estimated macro-model will generate stochastic values for certain

"output" variables, which can in turn be used to generate probability distributions for the second and any subsequent tiers of financial variables. And in principle it is clearly possible to incorporate further stochastic elements by making the intermediate relationships themselves stochastic although, for the time being at least, the Bank's approach is to treat the intermediate relationships as deterministic. The net result, in any event, is to yield a probability distribution for the value of some "target" variable, such as the profits or capital of the banking system.

14 To make this a little more concrete, if we were looking, for example, at the modelling of a shock to banks' corporate exposures, the process would start with variables such as GDP growth, short and long-term interest rates, inflation, unemployment and exchange rates as they are emerge from a macro-model and might then derive measures of the impact on corporate income, interest payments and collateral values. From these variables in turn a measure of the corporate default rate might be calculated, based on an historic reduced-form relationship between the "driver" variables and defaults. And from that the balance sheet impact can be determined by bringing in an estimate of loss-given-default. **[Slide 2.]**

15 Typically such models will be used to explore departures from a base case and therefore to estimate the impact of a particular scenario on some "significant" financial variables. Ideally, one of these variables would be a measure of financial stability. As I noted earlier, however, there is no fully satisfactory metric of this kind, certainly in terms of overall welfare loss, but some variables may be useful stand-ins – and at present in the Bank we have adopted the net worth of the major UK-based banks as a helpful summary variable.

16 The augmented macro-model approach has the merits of being relatively simple and of relating closely to accessible data. But it also has some significant weaknesses. First, it is essentially a reduced form rather than a structural approach; it is not rooted very obviously in behavioural relationships. Second, it typically makes no explicit provision for responses to financial stress by either companies or households as borrowers, or banks as lenders. Third, it does not capture the internal structure of, or interactions *within,* the financial sector, which arguably lie at the centre of any analysis of financial crises. For the same reason it has little or nothing

to say about contagion or crisis dynamics. Despite these weaknesses, the augmented macro-model approach has proved to be of considerable value as a basis for a variety of practical approaches to measuring and managing risk. It has been used, in particular, in the context of stress testing where it provides a way of constructing economically consistent scenarios and estimating their implications for the financial sector – or typically, but more narrowly, the banking sector.

17 The design and application of stress tests is an issue in itself which I am not going to pursue here. For firms, their use is now widely recognised as a valuable, indeed essential, complement to other more backward-looking techniques of risk analysis and management. Let me make just two points. First, the IMF now regularly include stress testing in their Financial Sector Assessment Programs (FSAPs) for individual countries, as a way of assessing the overall robustness of a country's banking sector. Many of the techniques were pioneered in the UK FSAP conducted in 2002. Second, stress tests derived from augmented macro-models are a key element in the more quantified approach to risk assessments which the Bank is now adopting in its *Financial Stability Report*. Although the margins of uncertainty are large, in terms both of the probability that a particular scenario develops and the impact if it does, this nevertheless represents a useful practical step towards the goal of setting financial stability analysis in a more rigorous quantified framework.

Dynamic stochastic general equilibrium models

18 As I have noted, a shortcoming of the augmented macro-model approach is that the financial implications appear as a "reduced form" add-on without clear behavioural underpinning and without much in the way of feedback to behaviour either in the financial sector or in other parts of the economy. Moreover, the lack of any internal structure in the financial system - to the extent that they figure at all, financial intermediaries are effectively regarded either as identical or as consolidated into a single entity - means that these models are incapable of capturing some of the most important characteristics of financial instability.

19 A good deal of work has been undertaken recently in an effort to remedy these defects. I am not going to attempt a proper review, but simply pick out as a good

example the papers produced by Dimitrios Tsomocos in conjunction with Charles Goodhart and a number of other collaborators. Under the catchy label of dynamic stochastic general equilibrium models, this work aims to set the exploration of financial stability in a rigorous general equilibrium framework incorporating some important structural features of real world economies and financial systems and optimising behaviour at the level of individual economic agents. In particular the models incorporate:

(i) the possibility of default, which as an example of discontinuous behaviour presents a technical challenge;

(ii) incomplete markets, which means that not all risks can be hedged (and which in turn, as a market failure ie something which private markets do not internalise, provides part of the rationale for financial regulation); and

(iii) perhaps most important, incorporates a non-trivial degree of internal structure in the financial system and allows for heterogeneous financial intermediaries, which then allows for behavioural interaction between those intermediaries.

20 The price which has to be paid for the introduction of this degree of micro-structure is that the models become very complicated and potentially intractable both in terms of estimation and in terms of simulating real-world scenarios. Where the balance is struck along the complexity/faithful-to-reality spectrum is as always a matter of judgment. And indeed for different purposes it may make sense to strike the balance in different places.

21 A recent attempt to operationalise these models has, however, been made by Aspachs, Goodhart, Tsomocos et al (2006). They argue that, ultimately, the effects of shocks on system stability can be represented by two factors – banking profitability and the probability of default. Accordingly, an index embodying these two factors could serve as a useful indicator of financial fragility. In their empirical work, they find that both profitability and probability of default must deteriorate beyond some threshold value before there is an impact on economic welfare. [**Slide 3]** indicates what such a measure for the UK might look like. Higher values for the index point to

financial instability. The measure suggests that the UK financial system has been stable for much of the past decade, with notable blips before that in the early 1990s (during the so-called "small banks crisis") and then in early 2000 (broadly contemporaneous with the dot.com bubble).

Network models

22 A third and, at least in my view, promising approach – for reasons I will spell out in a moment – is what I have called "network" models. These also seek to import some internal structure into the financial sector but they typically do it in a stylised way rather than through the micro-modelling of behaviour. In essence, they represent the financial sector – probably best thought of as the banking sector – as a set of nodes interlinked by a network of counterparty exposures. Each node (bank) is endowed with a certain amount of capital and has a very simple balance sheet comprising, say, holdings of some "outside" asset (bank loans) and interbank claims, and liabilities (deposits) to non-financial agents as well as other banks. The model can be elaborated to include additional features, for example a tradeable asset whose price is determined by (flow) supply and demand. In addition, rules are needed to specify what happens at each node in response to changes in the balance sheet, and in particular what happens if the net worth of a bank falls to zero, or becomes negative.

23 In most cases, such models do not allow "closed solutions" – or the assumptions which are needed to permit closed solutions are so restrictive as to undermine the value of the model. But they can be used for numerical simulations. Moreover, they allow a good deal of freedom in exploring the effects of different reaction rules and patterns of interconnection; and although they are difficult to parameterise in a way which captures any real world financial system, they may have useful insights to offer on the way in which structural features – the distribution of bank sizes, the number of counterparties a bank has, and so on – affect the way in which shocks are transmitted or amplified. Furthermore, and the reason I say they seem particularly interesting, discontinuities (for example bankruptcy) and other frictions (for example the impact of "fire sales" on asset values) at the level of the individual bank generate in a very natural way non-linear dynamics for the system as a whole - which is exactly what experience suggests does happen in financial systems

under serious stress. [**Slide 4.]** This sort of chaos theory approach seems worth pursuing further as an additional route into the analysis of financial stability.

24 Network models suggest that financial systems may be, at the same time, "robust yet fragile". Whilst greater connectivity amongst the nodes makes for wider distribution of risks and lowers the initial probability of individual default, the impact – should a crisis develop – could be significantly increased. Greater interconnectedness means that an institution which survives the impact effects of a shock is likely to be exposed to a greater number of defaulting counterparties in the second round. Such models also suggest that shocks which are indistinguishable *a priori* can have significantly different consequences. A system may experience 1000 shocks of roughly comparable size and remain resilient. But the 1001st shock, if it impacts on a particular structural weakness or pressure point, could have a dramatically different effect. As in so many other areas of financial analysis, the past may not be a good guide to the future!

The future

25 This brief review of alternative approaches to the quantification of systemic risk indicates that many issues remain to be explored. Certainly in the Bank we envisage a program of further work extending well into the future. The objective will be to develop an integrated framework within which to set the Bank's macro-prudential analysis and commentary. As with the more familiar models applied in the monetary area, such a framework for looking at systemic risk should help to ensure consistency in policy judgements and give us a better way of ranking what can otherwise become simply a "laundry list" of disparate risks.

26 The approach being developed in the Bank by my colleagues Andy Haldane, Prasanna Gai, Mathias Drehmann and others will aim to bring together structural models of credit and interest rate risk, to take into account also the amplification effects associated with liquidity risk on both the asset side and liability sides of banks' balance sheets and then integrate all this with a network model of interbank linkages so generating forward-looking *aggregate* loss distributions for the UK banking sector. In that sense it represents a fusion of the three approaches I have discussed.

[**Slide 5.]** But some parts of this programme, notably the modelling of liquidity risk, are very much in their infancy. The latter needs to reflect the intuition that "fire sales" by a firm in difficulty can push prices below fundamental value, to a degree which reflects the liquidity or otherwise of a financial market, and at the same time represent the time profile of prices following distressed sales. It will also be important to capture interactions between the asset and liability sides of the balance sheet. The non-linearities which these interactions generate will almost certainly fatten the tails on the distributions of system-wide losses by even more than implied by pure network effects. The cliff faces in my chart will be steeper, with the implication that the potential violence of crises is greater. To the extent that these loss distributions capture feedback and amplification effects as well as network interactions they might, however, begin to provide rough approximations to the welfare cost of financial instability.

27 The Bank is not alone amongst central banks in seeking to develop systemic risk models but its approach differs in a number of important respects. First, it explicitly recognises interest rate risk as well as credit risk in banks' balance sheets. Second, a more rigorous, macro credit risk model is used to characterise shocks, allowing policy makers to explore specific scenarios as well as stochastic simulations. And third, as I have noted, the approach aims to capture internal frictions in the financial sector as reflected in liquidity effects. In addition, and in parallel with the assembly of individual components, it also aims to incorporate, as a form of cross-check, asset valuations derived from financial market prices using Merton-type techniques.

28 An integrated "suite" of models of this kind, while it represents a considerable advance on the starting point, nevertheless leaves some important issues unresolved. Perhaps most important, the "true" costs associated with systemic risk – as I noted before, the additional costs beyond those which would arise anyway in the face of some external shock – need to be isolated and fed back in a consistent way into the macroeconomy. Only then will it be possible to calibrate systemic risk properly in terms of the impact on GDP. This is the province of the second group of models I referred to earlier. A second challenge is that episodes of severe stress are quite rare. Quantifying and calibrating the interactions in these circumstances between balance

sheets, asset prices and other variables is therefore difficult, although there may be useful insights from recent work on search costs relating to, for example, the pricing of corporate bonds. And third, there are, more generally, significant shortcomings in the data relevant to systemic risk analysis. Specifically, balance sheet data is often hard to come by, particularly for LCFIs, and is often out of date. The Merton approach helps to avoid these problems by using more timely market information; but it has important limitations, notably its assumption of normal distributions. In addition, our map of interlinkages in the financial sector is derived mainly from banks' large exposures data, which gives at best an incomplete picture especially when risk transfer has significantly widened the range of relevant institutions.

Conclusions

29 This has been a very rapid tour of some rather difficult territory but I hope it has given a flavour of one aspect of the Bank's financial stability work. It is certainly at the theoretical end of the spectrum but the motives, as I indicated at the beginning, are severely practical.

30 It would therefore be reasonable to ask how much practical impact the analysis to date has had. Those of you who are regular readers of the Bank's *Financial Stability Report* will have noticed a significant change of style and approach in the edition published in mid-2006. That included, for the first time, a partly qualitative, partly quantitative indication of the probabilities and potential impact of various kinds of shock. We will be updating that analysis in the next edition due to be published in a month's time. Even the enthusiasts would not claim great precision for the numbers and in any case, to borrow some vocabulary from mechanics, they are probably as much measures of stresses as of strains.

31 But looking further ahead, a framework does now seem to be emerging within which a more systematic discussion of financial risks can take place. In time – as we saw with macro-economic models – better quantification should be possible. But don't hold your breath. Alongside all of the hurdles I have mentioned there is one further difficulty – namely that the structure of the financial sector is changing and

arguably changing faster than that of any other sector in the economy. A substantial element of judgment, experience and plain hunch is, I am afraid, likely to remain in the business of handling financial crises.

32 A final point. Much of what I have talked about this evening has been the work of a group in the Bank headed by Andy Haldane with Prasanna Gai and Mathias Drehmann. I would like to thank all three of them for keeping me, I hope, more or less on the right track in these remarks.

ENDS