Systemic capital requirements

Summary of Working Paper no. 436 Lewis Webber and Matthew Willison

Banking regulation has historically focused on making a detailed assessment of risk at the level of individual banks' balance sheets. But it is possible that, in an interconnected system, banks that appear sufficiently healthy when viewed individually may collectively present a material threat to the solvency of the system as a whole. First, there may be similarities between banks' asset exposures that generate a tendency for banks' solvency positions to deteriorate and improve together. This can leave the system vulnerable to common shocks to the macroeconomy or to capital markets. Second, losses at an individual bank that are sufficient to cause it to default may trigger contagious failures of other banks in the system if they have extended it loans. Such contagious failures could trigger further rounds of contagious defaults in the banking system. System-wide losses could then far exceed the size of the initial shock.

Vulnerabilities of the system as a whole that cannot be identified by focusing narrowly on the health of individual banks suggest that a change in the way that risks to the banking system are assessed and prudential requirements for banks are calibrated could be beneficial. For example, capital requirements for banks could be set with the goal of achieving a level of systemic credit risk that a policymaker is willing to tolerate. This paper describes a system-wide risk management approach to deriving capital requirements for banks that reflect the impact their failure would have on the wider banking system and the likelihood of contagious losses occurring. These are referred to in this paper as 'systemic capital requirements'.

At the centre of the approach is the policymaker's optimisation problem. The policymaker is assumed to be interested in ensuring that the probability of banking system insolvency over a given time horizon is less than a chosen target level. This reflects the policymaker's systemic risk tolerance. The target could, of course, be achieved in all states of the world by setting very high systemic capital requirements. But the policymaker may also want to limit the potential inefficiency costs associated with regulatory capital requirements. If equity capital is more expensive than debt because of market frictions, higher capital requirements could, for example, increase the cost of bank lending to non-bank borrowers. The possible trade-off between financial stability and financial efficiency motivates a constrained optimisation problem, where a policymaker seeks to identify systemic capital requirements for individual banks that minimise the total level of capital in the banking system, subject to meeting their chosen systemic risk target. In other words, a policymaker sets banks' capital requirements to maximise efficiency subject to achieving a preferred level of stability. The solution of the constrained optimisation problem is a unique level of capital in the banking system and its distribution across banks.

Nested inside the policymaker's constrained optimisation problem is a simple structural model of a banking system in which shocks to banks' non-bank assets can cause insolvency. The underlying model further allows such shocks that originate outside the banking system to be transmitted and amplified through a network of interbank loans, so that credit losses spill over onto other banks when one or more banks become insolvent. The model captures two important drivers of systemic risk: (i) correlations between banks' assets (as a result of common exposures to non-banks), which may lead to multiple banks becoming insolvent simultaneously; and (ii) the potential for contagious bank defaults to occur because of losses on interbank lending.

The model is calibrated to resemble the major UK banks. It is used to illustrate how assessing risks only at the level of individual banks' balance sheets can lead a policymaker to underestimate the level of systemic risk in the banking system as a whole. The probability of very large losses crystallising in the banking system is greater when the potential for interbank contagion is taken into account, particularly when a number of banks have their balance sheet simultaneously weakened by losses on loans to non-banks.

The modelling choices in this paper reflect a trade-off between realism (complexity) and pragmatism (simplicity) in the description of credit risks facing an interconnected banking system. The paper uses a simplified description of the evolution of banks' balance sheets so that computational effort can be focused on solving the constrained optimisation problem faced by the systemic policymaker, taking into account the interlinkages between banks. As such, the primary focus of the paper is to obtain general insights into the properties of risk-based systemic capital requirements, rather than to calibrate precise nominal amounts that may be required to achieve particular risk targets in practice.

Systemic capital requirements for individual banks, determined as the solution to the policymaker's optimisation problem, depend on the structure of banks' balance sheets (including their obligations to other banks) and the extent to which banks' asset values tend to move together. Generally, banks' systemic capital requirements are found to be increasing in: balance sheet size relative to other banks in the system; interconnectedness; and, materially, contagious bankruptcy costs.

The paper illustrates, however, that risk-based systemic capital requirements would decrease during economic upswings and increase during downswings in tandem with measures of bank credit risk that are based on contemporaneous financial market prices, other things being equal. This procyclicality can be smoothed, to some extent, by using through-the-cycle measures of the riskiness of banks' assets. Nevertheless, the effect of such smoothing on the distribution of system credit losses is modest relative to the effect of cyclical changes to the composition of banks' balance sheets (leverage), suggesting a role for explicitly countercyclical capital requirements.

Estimating the impact of the volatility of shocks: a structural VAR approach

Summary of Working Paper no. 437 Haroon Mumtaz

A large body of empirical work has focused on estimating the impact of structural shocks on the economy. A large proportion of these studies employ vector autoregressions (VARs) — a system of equations where each variable depends on the lags of all variables included in the model. However, in their current form VAR models cannot directly incorporate the possible role played by the change in the volatility of the structural shocks as this is assumed not to have a direct affect on the variables included in the model. As shown in recent theoretical work, however, changes in shock volatility and uncertainty can have a direct impact on the macroeconomy. For example an increase in uncertainty may cause firms to pause hiring and investment decisions thus affecting real activity.

This paper proposed an extended VAR model which incorporates two additional features. First it allows the volatility of structural shocks to be time-varying. Second it allows for a direct impact of this time-varying volatility on the level of the variables included in the model. The paper describes an econometric method to estimate this extended VAR model.

We use the proposed model to estimate the possible impact of changes in the volatility of monetary policy shocks on the US economy. The monetary policy shock is identified from the data using two methods: (1) by assuming that these shocks have no impact on output growth and inflation for one quarter due to policy lags; and (2) by assuming that when these shocks lead to an increase in the federal funds rate this results in a contemporaneous reduction in output and inflation. In both cases, we estimate that the volatility of the monetary policy shock was high during the mid-1970s, the early 1980s and during the recent recession.

In order to gauge the impact of the volatility of the monetary policy shock, the model is simulated under the scenario where this volatility is assumed to double and no other shocks hit the economy. Under these assumptions, this change in volatility is estimated to reduce US GDP growth by 0.2% and inflation by 0.3%. However, once the importance of this volatility shock is considered relative to other shocks hitting the economy, its contribution is found to be small. This suggests that, in relative terms, changes in the volatility of monetary policy shocks are not economically significant.

How do individual UK consumer prices behave?

Summary of Working Paper no. 438 Philip Bunn and Colin Ellis

It is important for monetary policy makers concerned with meeting an inflation target to consider how prices behave. Nominal rigidities imply that prices cannot freely adjust, and the degree of nominal rigidity in the economy will influence the short-term impact of monetary policy on real activity and hence the response of inflation. This paper uses a database of over 11 million price quotes to investigate how individual consumer prices behaved in the United Kingdom between 1996 and 2006. These are the microdata that underpin the monthly consumer prices index produced by the Office for National Statistics. This work enables us to establish the facts about how frequently consumer prices change and how much they change by when they do change, and it should help us to improve our understanding of the nature of the nominal rigidities that exist in the economy. The results also help to establish which theories of pricing behaviour most closely represent the way in which prices are set in the real world, or at least in the UK economy.

This paper is the first to examine how UK consumer prices behave using the individual price quotes underlying the published aggregate inflation measure that is targeted by the Bank of England. This paper complements similar work on producer prices, which examines how prices behave further up the supply chain, and a recent survey of how firms set prices that was carried out by the Bank.

We find that 19% of consumer prices change each month on average, although this falls to 15% if sales are excluded. There is little evidence to support the presence of downward nominal rigidities in product markets, since 40% of all consumer price changes are decreases. UK consumer prices appear to be slightly more flexible than in the euro area, but they are less flexible than in the United States.

Consumer goods prices change more frequently than those of services, as on average 24% of goods prices change each month, compared with only 9% of services prices. At the component level, the prices of energy goods change the most frequently. The main service sector components all display a similar degree of price stickiness.

The share of prices changing each month varies across different years of our sample. There is some correlation between the share of prices increasing and the aggregate consumer price inflation rate. There are also some seasonal effects: prices are most likely to change in January and April and least likely to change in November and December. For consumer goods prices, the probability of a price change is highest in the month immediately following the previous change. As more time passes since the last price change, the probability of a price changing in any given month declines. For services, prices are most likely to change a year after the previous change, suggestive of annual price reviews. The probability of services prices changing in other months is broadly constant.

The distribution of the size of price changes is wide, although a significant number of changes are relatively small and close to zero. Around 60% of all price changes are between -10% and 10%, and the modal price change is an increase between 1% and 2%. The distribution of the size of consumer price changes narrows a little if sale prices are excluded. There are more small increases in prices and fewer price cuts for services than there are for goods, but there are considerable differences in the shape of the distributions of price changes at the component level. Prices that change more frequently tend to do so by less. This relationship appears to be particularly strong for services prices, but it also hold for goods prices as well once the effects of sales are taken out.

Our results on the behaviour of UK consumer goods prices are similar to those from previous work on UK producer prices (which covers only goods and not services). This suggests that there are few pricing frictions between the production and retail sectors in the United Kingdom.

Our findings from the microdata are not consistent with any one theory of price-setting. The marked heterogeneity that we observe in the behaviour of prices in different parts of the economy suggests that different theoretical models may better explain how prices are determined in different sectors. This would argue against the use of 'representative agent' models. The challenge is to develop a new theory of price-setting that better fits the stylised facts observed in these micro-studies while also fitting the properties of the aggregate macrodata.

An efficient minimum distance estimator for DSGE models

Summary of Working Paper no. 439 Konstantinos Theodoridis

Economic models are useful to economists and policymakers only if they are able to reproduce important features of the observed data. This property depends crucially on the values attached to model's parameters, and one way to decide about them is through the 'estimation' of the model. In essence, estimation is a mathematical procedure where the chosen parameter values minimise an objective function. A well-known example is 'least squares', minimising the squared distance between the actual data and the predicted values, which penalises large mistakes. Unfortunately, the estimation of modern macroeconomic models that rely heavily on microeconomic theory to explain the behaviour of economic agents and therefore the evolution of the economy over time while subject to random (stochastic) shocks (known as dynamic stochastic general equilibrium (DSGE) models) poses serious difficulties. This is due to the fact that theory imposes on the data a large number of very severe restrictions, which are not always supported by the latter.

Despite this, DSGE models are very useful. They are an abstraction of the economy that allows economists and policymakers to think clearly about economic relationships and actual developments, combining theory and data in a coherent way, and thus offering real insights. The way to make this work is to keep the model simple, meaning that a large number of strong restrictions need to be imposed on the data. This trade-off between the usefulness of the model and its ability to replicate elements of the true world is what makes the estimation of microeconomic theory founded models a challenging task.

The objective function used for the estimation of the model can be based on all available data information (full information) or on a few selected features of it (limited information). Full information sounds ideal, but in practice it makes large demands on the model. In the second case, the estimated parameters are chosen to minimise some measure of the distance between key characteristics of the data produced by the model and those observed in the data. One important feature that reveals the dynamic properties of the model is the 'impulse response function'. This shows the effect over time on a variable — say, inflation — after a shock hits the economy. (Indeed, many economists choose the parameters of their models judgementally in order to match the cyclical patterns of the data as they are summarised by the impulse response function — a process not of estimation but 'calibration'.) An advantage is that the targets that the estimated model aims to 'hit' are observed, meaning that failures to match these statistics of interest can be used to infer what parts of the theory are still missing from the model and derive useful economic conclusions. This is not true for full-information techniques where the estimated parameter vector minimises the distance between the model and the true data-generation process, which is unknown and highly abstract.

At the heart of the problem is that we cannot hope to explain everything in economics. A particular DSGE model is usually developed to explain only certain economic phenomena. Limited information estimation techniques let the model reproduce these facts as closely as possible. This increases the usefulness of the model since the user can immediately assesses how well the model serves its purposes of creation and, consequently, to decide whether it can be used to draw meaningful economic conclusions.

This study introduces an impulse response matching estimator that encompasses all the existing ones. It relies on the maximum information set (it mimics full-information estimators under some conditions), while existing methods utilise only a small part of the available set of instruments. The statistical theory (assuming we have a very large sample) developed here covers all the existing impulse response matching estimators and thus closes an important gap in the literature. The (more realistic) small-sample behaviour is investigated through a simulation exercise, where the proposed estimator is compared to other (modern and less modern) estimators for theory-driven models.

The measure that results from the estimation of the model can be used to assess whether a model's dynamic properties (as they are summarised by the impulse response functions) are statistically different from those observed in the real world, meaning that it can serve as a device to rank candidate economic theories that aim to explain the same features of the data. The work in this paper uses a widely used macroeconomic model to assess the usefulness of the method. The results are very promising. Now that the proof of concept has been established, the next step will be to apply the method to real, rather than simulated, data.

Time-varying volatility, precautionary saving and monetary policy

Summary of Working Paper no. 440 Michael Hatcher

In order to design effective monetary policy, central banks require an understanding of the mechanism by which economic shocks are transmitted to key macro variables like inflation, consumption and output. Economists therefore conduct policy analyses using models in which key economic relationships are spelt out but are subject to 'stochastic shocks' that represent unpredictable external events that influence the economy. A key task for monetary policy is to understand the transmission mechanism of such shocks, thereby enabling effective policy responses to be formulated.

Perhaps oddly, most policy analyses are carried out in a way that sidesteps the impact of uncertainty on households. Such models can match many features in the data and have a number of advantages. Notably, they can be represented in the form of a linear system of equations, making numerical simulations of medium and large-scale models feasible. However, an important drawback is that they cannot properly capture swings in uncertainty (fluctuations in the volatilities of economic disturbances), to understand the impact of such swings on the economy, or to evaluate potential policy responses. Yet, as exemplified by the recent financial crisis, changing uncertainty can be an important driver of economic behaviour. By ignoring such effects, these models provide policymakers with an incomplete picture and may lead to biased policy recommendations. Previous research at the Bank of England and elsewhere has examined the impact of uncertainty. But beyond that, there is an issue of whether changing levels of volatility also affect behaviour materially. This paper builds on that work and investigates the issue in more detail, focusing on a single aspect of household behaviour that is influenced by changes in uncertainty — precautionary saving.

Precautionary saving is additional saving driven by the possibility that if households are unlucky, consumption will fall to a low level, at which point an extra pound of spending is highly valued. This introduces a powerful non-linearity into economic models which has to be addressed explicitly. Furthermore, it has direct relevance for monetary policy, because an increase implies a reduction in current consumption, the main component of aggregate demand and an important factor influencing the extent of inflationary pressure in the economy. Thus we look at the monetary policy implications of ignoring precautionary savings effects arising from variations in the volatilities of demand and supply disturbances hitting the economy — an investigation which, by definition, cannot be conducted within a constant volatility framework.

In order to capture these effects in the model solution, the model is solved numerically using a higher-order approximation method. Given that the mechanism is driven by uncertainty, crucial to financial markets, consumer preferences are specified in a way that has been shown to provide a better 'match' to asset pricing data. Specifically, it is assumed that utility follows an 'external habits' specification, such that consumers value the difference between consumption and a slow-moving reference value. This specification of preferences introduces cyclical variation in risk appetite and raises household aversion to risk, two effects that appear to be important features of financial markets. Given that the model itself is stylised, the quantitative results reported are intended to illustrate rather than estimate the monetary policy implications of volatility fluctuations.

A key finding is that volatility fluctuations can have a small but relevant impact on precautionary saving behaviour, and therefore upon the appropriate conduct of monetary policy. The main contribution of the paper is to clarify the mechanism by which volatility fluctuations are transmitted through the precautionary savings channel and to illustrate — both analytically and quantitatively — the implications for monetary policy. If volatility fluctuations are not taken into account by policy, interest rates will be set incorrectly. As a result, a central bank that follows an interest rate rule that ignores volatility fluctuations will increase inflation and output instability, albeit to a small degree. Moreover, sensitivity analysis shows that the extent of 'policy bias' falls as the importance of habits in preferences is decreased. Consequently, models which are not calibrated to match higher-order risk effects may understate the importance of volatility fluctuations for the economy.