Reputation, risk-taking and macroprudential policy

Summary of Working Paper No. 462 David Aikman, Benjamin Nelson and Misa Tanaka

This paper considers the role of macroprudential countercyclical capital adequacy regulation in moderating credit cycles in a simple theoretical model. In our model, banks not only care about returns on their investments, but also their reputations. Imperfect information about banks' abilities and profitability means that they suffer a bigger reputational loss if they fail to make money when macroeconomic fundamentals are good than when they are bad. This is because when fundamentals are good, high-ability banks are more likely to earn high profits, such that markets attribute low profits to the low ability of bank managers. The fear of getting a bad market reputation gives low-ability bank managers the incentive to hide low profits and extend excessive credit in a bid to 'gamble for reputation' when fundamentals are good. This generates socially inefficient credit booms which ultimately lead to bank losses.

Our analysis suggests that countercyclical capital adequacy requirements are constrained socially optimal when macroeconomic fundamentals are within an intermediate range. By helping to reduce the incidence of inefficient credit booms, countercyclical capital adequacy requirements help to meet the dual objectives of moderating credit cycles and enhancing banking sector resilience. We are also able to separate two effects of countercyclical capital requirements on banks' risk-taking incentives, namely (i) the direct effect of raising the cost of risk-taking and (ii) the indirect effect of making information about the state of macroeconomic fundamentals public. We demonstrate that the latter can have a powerful effect in reducing banks' risk-taking incentives when fundamentals are rapidly deteriorating.

Our analysis focuses on a particular role for capital adequacy requirements, namely, that of preventing banks from investing in risky projects that have negative net present value. There are other rationales for countercyclical capital adequacy requirements which we have not considered here, including enhancing loss absorbance and avoiding socially costly financial crises. Our analysis also focuses on the role of capital adequacy requirements in preventing inefficient credit booms, and does not examine its potential role in preventing inefficient credit crunches. Examining all these aspects of countercyclical capital requirements in a single framework is left for future research.

The international transmission of volatility shocks: an empirical analysis

Summary of Working Paper No. 463 Haroon Mumtaz and Konstantinos Theodoridis

The recent financial crisis has been characterised by increasingly volatile macroeconomic data in the United States and the United Kingdom. In this paper we devise an empirical model to estimate the impact of this increase in volatility or uncertainty on the UK economy. In particular we examine the impact of an increase in uncertainty associated with US real activity. Uncertainty about growth in large economies has been a key consideration for policymakers in recent years.

The empirical model that we propose is an extension of vector autoregression (VAR) models. VAR models link each variable included in the model to past values of all the variables in the system. The residual associated with each variable is typically assumed to have a constant variance. For example if the model included US GDP growth, the variance of the residual to the relevant equation would be constant. This also implies that in this modelling set-up, the uncertainty associated with each variable (as proxied by the residual variance) is fixed over time. Given recent events, this may not be a good assumption.

Our paper extends this model along two dimensions. First, we allow the residual variance to change over time — in other words we allow for stochastic volatility. Second, we allow this stochastic volatility to enter as an explanatory variable in each equation of the model. We can therefore gauge the effect of volatility on each variable included in the VAR model.

In our empirical application, we include US GDP growth, US CPI inflation, the federal funds rate, UK GDP growth, UK CPI inflation and Bank Rate in the extended VAR model. We then try to estimate the impact of an increase in the stochastic volatility associated with the residual of the US GDP growth equation. We find that if this volatility increases by one standard deviation, UK GDP growth declines by 0.1% and UK CPI inflation increases by 0.1%. The impact of this shock on the US GDP growth and inflation is very similar. The impact is statistically important albeit small in economic terms.

We then employ a theoretical model of the open economy to understand the transmission channel of this shock. Model simulations indicate that it can be interpreted as a sudden change in the volatility associated with shocks to US wages or productivity — ie shocks to US 'supply'. A sudden increase in the volatilty of these shocks leads to an increase in precautionary savings by consumers who are more uncertain about the future. This leads to a reduction in consumption and subsequently GDP growth in both countries. Workers try to insure themselves against uncertainty about future wages by demanding higher pay in the current period and this puts upward pressure on inflation.

International policy spillovers at the zero lower bound

Summary of Working Paper No. 464 Alex Haberis and Anna Lipińska

In this paper, we are interested in how the policy of other central banks affects policy in a small open economy in the face of a large global demand shock that leads central banks internationally to cut rates to the zero lower bound (ZLB). Our interest in this issue comes from the policy response to the financial crisis that started in 2007/08. This hit many economies at the same time, leading to large declines in output during what has become known as the 'Great Recession'. In response, central banks around the world cut policy rates to (close to) zero to offset the deflationary pressure associated with the collapse in demand.

The ZLB creates an interesting set of challenges for monetary policy. This is because the conventional options available to policymakers to stimulate the economy — further rate cuts — are not available. Past academic work has shown that this can lead to trade-offs for policymakers in terms of stabilising inflation and output. In the current conjuncture, with the crisis having led many of the world's major central banks to cut policy rates to (or close to) the ZLB, the international dimension of these challenges is of particular interest — for instance, in terms of how policy overseas might create spillovers into the policy problem at home, which is the focus of this paper.

In practice, however, it is worth noting, that central banks greeted these challenges during the crisis with 'unconventional' quantitative policy action. In this paper, we do not look at unconventional monetary policy measures taken at the ZLB.

Instead, we adopt a more stylised framework for looking at monetary policy strategy, in line with previous research on monetary policy at the ZLB. We adopt this approach in part for its analytical convenience and in part because it allows us to couch our findings in terms of other work. In this framework, policy may be set either under 'discretion' or under 'commitment'. Discretionary policy involves the policymaker taking the action in a given period that gives the best outcome in terms of stabilising inflation and output in *that* period. When following policy under commitment, the policymaker commits to the course of action for all periods that achieves the best stabilisation performance over time.

This is more powerful, because, if it is possible, policy can operate more effectively on expectations about the future, which under discretion are constrained by the belief that policymakers will choose short-sighted policies (the famous 'time consistency' problem). In general, at the zero bound, commitment policy allows the policymaker to provide greater stimulus to the economy, which leads to improved stabilisation of inflation and output relative to a purely discretionary policy.

To analyse the issue, we use a model in which there are two countries: a large economy (which we refer to as 'foreign') and a small open economy (which we refer to as 'home'). The foreign economy is large in the sense it is not affected by developments in the home economy, although developments in the foreign economy can affect the home economy.

In our results, we find that in response to a large global demand shock, when foreign policy follows a commitment strategy, this reduces the home policymaker's ability to stabilise the home economy when home and foreign goods are close substitutes. This is because looser monetary policy in the foreign economy means the home real exchange rate is relatively appreciated compared to when the foreign policymaker sets policy under discretion. When there is a high degree of substitutability between goods, a stronger home real exchange rate induces large expenditure-switching effects away from home goods. This effect outweighs the impact on the demand for home goods from the higher level of foreign aggregate demand resulting from the looser stance of foreign monetary policy. Because our model is based on microeconomic foundations, we are able to work out how foreign policy affects social welfare at home. When goods are highly substitutable, home welfare is higher when foreign policy is set under discretion compared to commitment. In contrast, when goods are not close substitutes across countries, the opposite holds.

Size and complexity in model financial systems

Summary of Working Paper No. 465 Nimalan Arinaminpathy, Sujit Kapadia and Robert May

What role do large banks play in systemic risk and financial instability? How should capital adequacy standards recognise this role? How is stability shaped by concentration and diversification in the financial system? This paper explores these questions using a deliberately simplified, dynamical model of a banking system.

Developing methods used in epidemiology and ecology, we adopt network techniques which are well suited for such questions, particularly in modelling 'contagion' that is transmitted through linkages in the financial system. Specifically, we bring together three important transmission channels into a unified framework: (i) liquidity hoarding, where banks may cut their lending to each other as a defensive measure; (ii) asset price contagion linked to the falls in market prices which may be generated by asset sales by banks in distress; and (iii) the propagation of losses which may occur if banks default on their obligations to other banks in the interbank market (the network of lending exposures among banks). Importantly, we also integrate a mechanism for capturing how broader swings in 'confidence' in the system may contribute to instability, with the overall state of the system potentially influencing an individual bank's actions, and *vice versa*.

The interaction of such network and confidence effects arguably played a major role in the collapse of the interbank market and global liquidity 'freeze' that occurred during the financial crisis. Interbank loans have a range of maturities, from overnight to a matter of years, and may often be renewed, or 'rolled over', at the point of maturity. A pronounced feature of the 2007–08 crisis was that, as the system deteriorated, banks stopped lending to each other at all but the shortest maturities. The bankruptcy of Lehman Brothers in September 2008 transmitted distress further across the financial network. The effects extended well beyond those institutions directly exposed to Lehman Brothers, with banks throughout the system withdrawing interbank lending outright and propagating distress to the real economy by sharply contracting household and corporate lending.

Several specific motivating factors have been proposed to explain 'liquidity hoarding' (the maturity-shortening and ultimate withdrawal of interbank lending): precautionary measures by lending banks in anticipation of future liquidity shortfalls; counterparty concerns over specific borrowing banks; or collapses in overall system confidence. Our framework parsimoniously incorporates all of these mechanisms, while also capturing the idea that a bank's distress may affect not just those directly exposed or linked to it, but also confidence in the market at large.

We use our model to explore the effects of shocks to the system, such as the failure of banks or big losses on certain types of lending. We focus particularly on the adverse feedback dynamics arising from each of the contagion channels included, the effects of size disparity among banks and system concentration, and the effects of diversification. Our results highlight the disproportionate importance of large, well-connected banks for system stability: the impact of their collapse arises not only from their connectivity, but also from their effect on confidence in the system. Moreover, we show that while diversification may serve to limit the risk of failure of an individual bank, it does not mitigate the importance of that bank to systemic risk, and may indeed exacerbate it. Overall, these results illustrate the different approaches needed for regulation focused at the level of individual banks, and that focused on a systemic level. While sound microprudential regulation remains important for the former, the latter, macroprudential perspective, supports the notion of regulatory requirements concomitant with bank size, interconnectedness or (more generally) systemic importance. In particular, imposing tougher capital requirements on larger banks than smaller ones can enhance the resilience of the system. Furthermore, such requirements may also have the beneficial side-effect of providing disincentives for financial institutions to become 'too big to fail'. Our findings have conceptual analogies in ecosystem stability, and in the control of infectious diseases, which we also discuss briefly.

As with any theoretical approach, there are important caveats to our model. In particular, a key empirical challenge for future work is to quantify the confidence processes which we model. Incorporating uncertainty, for example over the underlying health of individual institutions or the system as whole, would also be a useful extension. Another key question is how the vulnerabilities in financial systems modelled in this paper emerge, and potentially grow, over time. Finally, while this paper focuses on one aspect of the regulatory response relating to capital requirements, other policy responses, such as the use of liquidity requirements or the implementation of effective resolution regimes, are also likely to be important in enhancing the resilience of the financial system.

QE and the gilt market: a disaggregated analysis

Summary of Working Paper No. 466 Martin Daines, Michael A S Joyce and Matthew Tong

In response to the deepening financial crisis in Autumn 2008, central banks in advanced economies reduced their policy rates sharply and introduced a range of other more or less unconventional measures designed to ease monetary conditions and to support financial stability.

In the United Kingdom, a key element of these unconventional monetary policy measures has been the programme of asset purchases financed by central bank money, commonly described as quantitative easing (QE). During the first round of QE purchases over the period from March 2009 to January 2010, the Bank of England (through an indemnified Asset Purchase Facility) bought £200 billion of domestic private and public assets, the vast majority of which were medium to long-term UK government bonds (gilts). By the end of January 2010, the Bank's gilt holdings represented nearly 30% of the stock of nominal gilts outside the official sector. Subsequently, between October 2011 and May 2011, the Bank completed a further £125 billion of purchases and, more recently, at its July 2012 meeting, the MPC voted to increase the size of its asset purchase programme by a further £50 billion to a total of £375 billion.

In this paper we analyse the impact on the gilt market of the first round of QE purchases during March 2009 to January 2010, in order to draw out lessons both about the effectiveness of the policy and also to shed light on the nature of the transmission mechanism from purchases onto bond prices/yields — a key link in the transmission of QE to the wider economy. In conditions where markets are functioning efficiently, one might expect economic news to be quickly assimilated into market prices as soon as it becomes available to market participants. But, given the unprecedented nature of the QE policy and market conditions at the beginning of 2009, it seems possible that the effects of QE may have taken longer than normal to get reflected in prices and indeed that the full market adjustment might have been delayed until the asset purchases were actually made. The contribution of this paper relative to earlier work on QE is to look at the effects of both the announcements (news) about QE and the actual purchases through the Bank's reverse auctions using disaggregated high-frequency data.

Our analysis of the high-frequency market reactions to individual announcements on QE suggests that the initial impact from the announcements took time to be fully priced in and that the cumulative initial impact on yields varied significantly across the term structure, with the largest impact up to 120 basis points between the 15 and 20-year maturity. We also find evidence that gilts with maturities close to or in the Bank's purchase range experienced larger relative yield falls (consistent with 'local supply effects') and that yields also fell more for gilts with longer maturities (consistent with 'duration risk effects').

Analysis of the Bank's reverse auctions suggests that ahead of each auction they led to further yield reductions on gilts both eligible and ineligible for purchase that averaged 2.5 basis points and 1.5 basis points respectively. These effects were not always reversed before close of business on the same day, with more persistent effects found to be positively associated with the degree of price dispersion of the accepted offers, an indicator of price uncertainty. These persistent effects may partly reflect learning by market participants. In addition, we find that the importance of the overall effects of the auctions on gilt yields diminished over time, as both liquidity and market functioning improved and knowledge of the operation of the Bank's purchase programme increased.

Econometric analysis of the time-series behaviour of gilt yields is consistent with the QE effect on gilt yields being quite persistent, once allowance is made for the countervailing effects on yields of fiscal news and improving macroeconomic prospects during 2009. Putting this finding together with our other results suggests that the peak gilt market response to the Bank's QE policy may not have occurred until the auction purchases began and the market learnt about the effects of the policy.

Overall our results suggest that the Bank's QE asset purchases had a significant and persistent impact on gilt yields. Our paper also provides direct evidence of local supply and duration risk effects consistent with imperfect substitution, which has implications beyond the financial crisis for how we think about price determination in the gilt market.

Factor adjustment costs: a structural investigation

Summary of Working Paper No. 467 Haroon Mumtaz and Francesco Zanetti

The costs associated with changes in capital and labour inputs are important factors affecting firms' decisions to expand or contract production. These ultimately affect the levels of economic activity and the patterns of business cycle activity that an economy experiences over long periods, and understanding the process is consequently important to macroeconomic policy makers. This paper investigates what theory and data tell us about the precise nature of adjustment costs, thus enabling us to build macroeconomic models better to describe business-cycle fluctuations.

We conduct the analysis by estimating a 'dynamic stochastic general equilibrium' model that accounts for several important features of the economy. Dynamic, because it emphasises how the economy evolves over time; stochastic, because in the model as in the world agents are continually buffeted by random shocks of various kind; and general equilibrium, because all parts of the economy are connected and affect each other. We examine several competing adjustment costs functions using US aggregate data. This approach has two main advantages. First, the model is derived by solving the optimal decision of each agent in the economy, thus enriching our theoretical understanding of how adjustment costs affect production. Second, rather than estimating adjustment costs functions using single equations, we pursue a multivariate approach by estimating the entire structural model, enabling more accurate estimates, aided by the fact that the independent variables are uncorrelated with the error terms (shocks) in the model.

We also find that the empirically acceptable adjustment costs function is non-linear, is increasing in both labour and capital, and also accounts for joint interactions between the two production inputs. Alternative specifications, with only capital or labour adjustment costs are not powerful.

We find that adjustment costs are small for both input factors. According to the theoretical framework, total adjustment costs represent 1.98% of total output per quarter. In addition, the cost of hiring an additional worker amounts to fourteen weeks of wages, whereas the cost of an extra unit of investment equals 0.21% of average output per unit of capital. Such estimates are within the range of values estimated using disaggregated data.

The analysis suggests that the reaction of factor adjustment costs to shocks is generally procyclical, except to shocks to the rate at which jobs and capital are dismissed. Finally, technology shocks are a major influence on fluctuations in factor adjustment costs in the short run, whereas shocks to the job dismissal rate compete with technology shocks to explain the bulk of fluctuations of factor adjustment costs in the long run.

Using Shapley's asymmetric power index to measure banks' contributions to systemic risk

Summary of Working Paper No. 468 Rodney J Garratt, Lewis Webber and Matthew Willison

Policymakers have in the period since the crisis been discussing how to regulate banks in ways that reflect the potentially different contributions banks make to systemic risk in the financial system in the event of their failure. One aspect of how an individual bank's failure could contribute to systemic risk could be defined in terms of whether its failure is considered to be pivotal in tipping the banking system from a state of stability to a state of instability. Based on this idea, we develop an approach that can be used to calculate the marginal contributions of individual bank failures to systemic risk.

The approach is based on a measure originally introduced by the mathematician and economist Lloyd Shapley. The so-called Shapley value is a way of allocating the output produced by a group among its members in a way that reflects fairly their individual contributions. In this paper we apply the Shapley value to the situation where the group is a set of banks that fail due to shocks to the values of their assets and the good they produce is in fact something bad — in this approach the bad is the failure of a set of banks tipping the system from a state of stability to one of instability.

The framework requires two key inputs: the values of banks' exposures to different asset classes; and the levels of banks' capital available to absorb losses on their asset holdings. The banking system can be hit by a range of shocks, which are defined in terms of the extent to which they reduce the value of the different asset classes. The shocks are assumed to occur with equal probability. For each possible shock, banks can be lined up in the order that they would fail as a result of that shock. Banks with asset portfolios weighted more towards the assets affected more by the shock, and/or have lower levels of capital, tend to be higher up the order of failure. The pivotal bank is the one that, when it is added to the banks that fail before it, causes the value of the failed banks' assets to move above a critical threshold value — this is defined as a systemic event. The pivotal bank receives a score of one (and other banks receive a score of zero). By taking an average of a bank's score over the range of possible shocks we calculate a measure of a bank's contribution to systemic risk. We illustrate, using simple examples, how banks' contributions depend on their asset portfolio compositions and their capital levels as well as on the calibration of the critical threshold that defines a systemic event.

We outline several ways in which the framework could be extended to consider: different definitions of a systemic event; adjustments to the values of banks' asset exposures to reflect the riskiness of those exposures; and the possibility of interbank contagion. We conclude by identifying some possible key next steps and further extensions of the approach. A key next step will be to apply the approach to bank data so that it could be used as a risk assessment tool. Since our approach applies to circumstances in which the system is in a state of instability, it would be natural to use our approach as part of a reverse stress-testing exercise.

High-frequency trading behaviour and its impact on market quality: evidence from the UK equity market

Summary of Working Paper No. 469 Evangelos Benos and Satchit Sagade

This paper studies the behaviour of high-frequency traders (HFTs) in the UK equity market and analyses its impact on aspects of market quality such as liquidity, price discovery and excess volatility. Although there is no precise definition of an 'HFT', the term is commonly used to describe firms that use computers to trade at high speeds and who also tend to end the day flat, ie carry small or no overnight positions.

HFT activity has increased steadily over the recent years in the US, the UK and continental European equity markets and, following a number of market mishaps (which seem to have been triggered by flawed computer trading algorithms), high-frequency trading has also caught the attention of regulators. However, the empirical evidence on the behaviour and impact of HFTs has so far been relatively limited and inconclusive. Thus, the Bank of England has a natural interest in better understanding HFT behaviour and how it might impact the quality of UK equity markets. In particular, a key question is whether and how HFT activity impacts price efficiency and liquidity.

This paper uses a sample from a data set of transaction reports, maintained by the Financial Services Authority, to attempt to give a first answer to these questions. The data identifies the counterparties to each transaction, which enables us to identify HFTs and study their behaviour.

We first find that HFTs exhibit substantial variability in their trading strategies. For instance, while some HFTs trade primarily passively (by posting orders that rest on the order book of the exchange so that others can trade against them), others trade primarily aggressively (by trading against resting orders of passive traders). In other words, some HFTs mostly supply liquidity and others mostly consume it. For this reason and in order to examine how trading behaviour is related to these patterns of liquidity provision, we split the HFTs in two groups, according to their trade aggressiveness, and examine the behaviour and impact of each group separately.

The 'passive' HFTs tend to alternate their positions over the short run (ie their buys tend to be followed by sells and their sells by buys) and their positions also tend to be insensitive to recent price changes. Conversely, 'aggressive' HFTs do not alternate their positions, and tend to trade in the direction of the recent price trend (ie they buy when the price rises and sell when it drops).

We next examine whether and how price volatility and the prevailing bid-ask spread influence HFT activity. We find that both 'passive' and 'aggressive' HFTs trade relatively more when prices are more volatile and when the spread is narrow. We suggest a number of reasons why this might be so.

Finally, we examine the impact of HFT activity on volatility. We note that volatility can be either 'good' (when price changes reflect the arrival of new information about fundamentals) or 'excessive' (when price changes do not reflect any information about fundamentals). In the latter case it is also referred to as 'noise'. Clearly, markets are more informationally efficient when there is more 'good' volatility and less 'noise'. We therefore examine the contribution of HFTs to both 'good' volatility and 'noise'. For that, we use an econometric framework that takes into account the exact time sequencing of HFT trades and price changes and, as such, allows us to isolate and estimate the causal effect of HFT activity on price volatility.

Our results show that there are instances where HFTs contribute (in absolute terms) a large amount of both 'good' and 'excessive' volatility; more so than the average, non-HFT, trader. This is possible if some of their trades carry a large amount of information while other trades are uninformative. We hypothesise that this may be because HFTs aim to end each day with relatively flat positions: if an HFT must, at some point during the day, only trade in order to adjust their inventory, these trades will have no information content and will likely create noise. For the stocks we analyse, HFTs are more informationally efficient than non-HFTs as their relative contribution of 'good' to 'excessive' volatility is on average 30% higher than that of non-HFTs. Owing however to the small number of stocks in our sample, we cannot confidently generalise these findings in the entire cross-section of stocks.

Given the instances of large contributions of both 'good' and 'excessive' volatility by HFTs, it is not immediately clear what the welfare implications of HFT activity are. If improvements in price efficiency at some times come part and parcel with additional noise at other times, then whether HFT activity is socially beneficial or not, will ultimately depend on how the marginal utility of information compares with the marginal disutility of noise, ie on how much additional noise we are willing to tolerate at some times for the benefit of more informed trading at other times. It will also depend on the balance between any beneficial impact HFTs may have on markets during 'normal' market conditions and the effect of HFT activity under more 'stressed' market conditions. Finally, the welfare implications of HFT activity will also depend on the propensity of errors in the operation of their algorithmic trading to cause harmful disturbances of the type experienced in the 'Flash Crash' of 6 May 2010. However, these issues are beyond the scope of this paper.