Bank behaviour and risks in CHAPS following the collapse of Lehman Brothers

Summary of Working Paper No. 451 Evangelos Benos, Rodney Garratt and Peter Zimmerman

During the period of financial stress, in the wake of the Lehman Brothers' default, infrastructures used by banks to make payments to one another held up well. The Bank of England's *Payment Systems Oversight Report 2008* explains that although the crisis placed unprecedented demands on payment and settlement systems, these continued to provide a robust service. We examine how this stress affected payment patterns in CHAPS, the United Kingdom's large-value wholesale payment system. This is important to the Bank in its role as the overseer of recognised interbank payment systems, including CHAPS, and as host of the infrastructure that supports the operations of CHAPS.

CHAPS payments data show that, in the two months following the failure of Lehman Brothers, banks on average made payments at a slower pace than prior to the failure. This delay was partly explained by concerns about bank-specific and system-wide risks. 'Turnover', which is defined as the average number of times each unit of liquidity employed by banks to make payments is used during the day, was 30% lower in the period from 15 September 2008 to 30 September 2009 than in the period preceding the Lehman default. In the immediate aftermath of Lehman this was due to payment delay, but later may have been related to increased reserves balances associated with quantitative easing. This may have led to banks being more willing to make payments with their own liquidity rather than relying on liquidity from payments received from others.

We also find that the payment delays observed in the months following the failure of Lehman Brothers modestly increased the liquidity risks associated with operational outages. An operational outage is an event during which a single settlement bank (ie a bank which is a member of CHAPS and is able to submit payments directly into the system) may be unable to send payments. Since such a settlement bank is unable to provide liquidity to the payment system, the impact of an operational outage depends on the liquidity that the affected bank would have provided to the system during the outage.

We compute two estimates of the impact of operational outages. One measure considers the impact of a single outage that occurs at the worst possible time on a given business day, while the other measure computes the expected impact of a single outage occurring at a random point in time during the day. Both measures of risk show a statistically significant increase in the period following the collapse of Lehman Brothers. Thus, our results show that, although operational risks did not crystallise, the potential for disruption in CHAPS did increase during the period of financial stress in the wake of the collapse of Lehman Brothers.

To provide some indication of the economic cost of these risks, we calculate how much additional money banks would on average have had to pay to insure themselves against the loss of liquidity due to an operational outage. Although the amount of liquidity loss to be insured against increased in the wake of the Lehman Brothers' collapse, a mitigating factor to this increase was a sharp decline in the cost of obtaining liquidity during the same period. The combined effect was an increase in the hypothetical premium until mid-October 2008, followed by a fall to levels lower than those seen in Summer 2008, on account of lower borrowing rates. Despite the temporary increase, the daily hypothetical premium was about £6,700 per bank during the month after the Lehman Brothers' collapse. While the economic cost was low, in absolute terms, an interesting question is whether the cost — and the underlying risk exposure — would have increased to a greater extent in the absence of CHAPS throughput requirements, which oblige settlement banks to settle minimum proportions of their payments by specific times of the day.

Simple banking: profitability and the yield curve

Summary of Working Paper No. 452 Piergiorgio Alessandri and Benjamin Nelson

This paper examines the relationship between bank profitability and interest rates. Understanding this link is important for policymakers. If interest rates have a systematic effect on bank profitability, and if in the short run profitability is a major determinant of bank capital, it follows that monetary policy may have implications for the resilience of the financial system. We investigate the effects of interest rates on profitability using a new, unique panel data set containing information on the UK activities of UK and foreign banking groups for 1992–2009. We find evidence of a systematic effect of market interest rates on bank profitability. In the long run, high yields and a steep yield curve boost banks' income margins. In the short run, though, an increase in short-term yields depresses income, which is consistent with the presence of frictions affecting the repricing of banks' assets and liabilities in an asymmetric way.

We begin with a simple theoretical model of a bank which is subject to credit and interest rate risk, which chooses its interest margin to maximise expected profits. The model provides us with a number of testable implications. First, in equilibrium the net interest margin (NIM) is likely to be positively related to short-term interest rates, as banks raise their loan rates and shrink their lending quantities in response to higher market rates. Second, the short-run and long-run effects of interest rates can differ. In particular, if banks borrow short and lend long, and if their interest rates are not fully flexible in the short run, banks will be exposed to 'repricing' risk. The combination of maturity mismatch and repricing frictions is indeed a popular explanation for why sharp changes in interest rates might compress bank profits. We find that high interest rates are associated with large interest income margins, as predicted by the model. We also find that the slope of the yield curve matters positively for interest income: after all, banks indeed seem to borrow short and lend long. The short-run impact of an increase in short-term market rates, however, is negative. This is consistent with the existence of significant repricing frictions that prevent banks from implementing their pricing decisions instantaneously. We also find that level and slope of the yield curve affect the net interest margin and trading income in opposite directions, which suggests that banks hedge interest rate risk through derivatives. Even after accounting for hedging, however, large banks appear to retain a residual exposure to UK interest rates: the interest rate effects in the banking book 'pass through' into operating profitability. Thus monetary policy — set for the economy as a whole — appears to have systematic effects on bank profitability, providing one potential motivation for the use of macroprudential policy tools.

We present two applications of our estimated model. First, we explore the interaction of level and slope effects and short and long-run multipliers by running a 'monetary policy shock' through the model. A typical policy tightening raises short-term rates and flattens the yield curve, thus depressing banks' income through two distinct channels. This effect is fairly short-lived, and somewhat attenuated by hedging. Higher rates have an unambiguously positive effect on bank profits in the long run. Second, we use our estimated NIM equation to decompose the sources of profitability since 1992, examining the model-implied contributions of the level and slope of the yield curve to the average net interest margin over the sample.

Estimating probability distributions of future asset prices: empirical transformations from option-implied risk-neutral to real-world density functions

Summary of Working Paper No. 455 Rupert de Vincent-Humphreys and Joseph Noss

There is a strong tradition of central banks and other policymakers extracting information from the prices of financial securities. Derivatives contracts can provide information on the expected future path of their underlying asset's price that goes beyond its central expectation. They therefore offer an insight into the level of uncertainty surrounding future cash flows. The Bank of England regularly estimates probability density functions (pdfs) from options prices in order to obtain an indication of the weight investors place on different future prices.

However, such option-implied pdfs may not provide a true indication of the actual probabilities investors ascribe to certain outcomes. This is because such pdfs give an indication only of the probabilities investors would have in mind if they were 'risk-neutral', and did not consider the uncertainty around an asset's future pay-offs in assessing its value. In the likely case that investors are averse to this risk, this would lead to differences between the risk-neutral densities backed out of options prices, and the true 'real-world' probability densities considered by investors.

The resulting estimated 'real-world' pdfs offer a number of advantages over their risk-neutral counterparts. First, they afford an insight into market participants' actual views on future asset prices, and offer an improved quantification of the uncertainty around financial variables. Second, a comparison of the risk-neutral and estimated real-world pdfs reveals new information as to how investors' risk preferences are affecting derivatives prices. Finally, estimated real-world probability densities are directly comparable with other forecasts considered by policymakers that are *not* based on derivatives prices, for example those of GDP growth and inflation.

The approach examined here is empirical in that it compares the risk-neutral distribution generated directly from options prices to the *actual* distribution of prices as they are later observed. To the extent that the two show a systematic disparity over time, this may be exploited to adjust the risk-neutral densities over as yet unobserved future prices to estimate the agents' real-world expectations.

This work offers a robust means of transforming risk-neutral densities obtained from options contracts on the FTSE 100 and short sterling. The resulting real-world probability densities offer a superior average fit across the distribution of observed prices than their risk-neutral counterparts. The resulting parameters appear stable over time, at least until the end of our data sample in June 2007. To the extent that this remained the case when the methodology was applied to prices since, it could form the basis of an operational method to better predict their future prices and enhance conjunctural analysis. It could also form the basis of more advanced work that aimed to condition this transformation on some other (macroeconomic) observable variable which may increase the method's predictive power.

Liquidity risk, cash-flow constraints and systemic feedbacks

Summary of Working Paper No. 456 Sujit Kapadia, Mathias Drehmann, John Elliott and Gabriel Sterne

Although the failure of a financial institution may reflect solvency concerns, it often manifests itself through a crystallisation of liquidity risk associated with a loss of funding. In such funding crises, the bank's solvency position no longer fully determines its survival; and its cash-flow constraint becomes critical.

This paper develops a framework that promotes an understanding of the triggers and system dynamics of liquidity risk during periods of financial instability and simulates the impact of these effects in a quantitative model of systemic risk. By using simple indicators and analysing bank-specific cash-flow constraints, we assess the onset and evolution of liquidity stress at individual institutions in various phases. And we capture several systemic feedbacks which may arise during funding liquidity crises, mostly linked to defensive actions taken by banks in distress, and many of which have been important during the current financial crisis. A key contribution of this paper is to demonstrate how systemic risk may escalate and contagion may spread to other institutions as a bank's funding conditions deteriorate, irrespective of whether the bank ultimately survives or fails. By applying the model to the UK banking system based on the balance sheet vulnerabilities that existed at the end of 2007, we illustrate how liquidity feedbacks may markedly amplify other sources of risk.

The severity of an individual bank's funding distress is calibrated using a simple 'danger zone' approach that scores each bank according to eight indicators that proxy solvency, liquidity profile, and confidence. Two indicators in particular play an important role in the transmission dynamics of funding crises modelled here. The first is short-term wholesale maturity mismatch (between contractually maturing liabilities and assets); a bank with a larger share of short-term borrowing faces greater funding liquidity risk. The second is that distress at one bank may adversely affect 'similar' banks through a pure confidence channel.

A danger zone score beyond a first threshold triggers the initial phase of distress, in which long-term unsecured funding markets close to the bank. The bank has to refinance a larger volume of liabilities in short-term markets each period, which further worsens its maturity mismatch score. The bank takes the defensive action of hoarding liquidity (reducing the maturity of its own intra-financial system lending), which improves its own mismatch score but worsens mismatch and increases danger zone scores at counterparty banks.

A second phase of distress is triggered beyond a further danger zone threshold. In the model, this results in shorter-term

unsecured funding markets closing to the distressed bank, which then takes further defensive actions in an attempt to meet its cash-flow constraint. If profits earned over the period are insufficient to meet liquidity needs, the bank in the first instance withdraws all maturing intra-financial system assets, using the proceeds to pay off liabilities due. Its next line of defence is to sell or encumber its liquid assets. Finally, it resorts to fire-selling its illiquid assets, precipitating falls in asset prices and generating systemic feedbacks as other banks holding those assets are assumed to suffer temporary losses, worsening their solvency position and potentially increasing their danger zone score. If the combined effect of these defensive actions is insufficient for the bank to meet its cash-flow constraint, it fails. At this point, it defaults on its obligations to other banks, with the associated counterparty credit losses determined using a network model of bilateral interbank exposures. In extreme circumstances, the spillover effects linked to liquidity hoarding, asset fire sales, confidence channels and counterparty default may also generate sufficient contagion to cause other banks to suffer funding liquidity crises, and potentially fail.

The paper provides illustrative simulations using a version of the Bank of England's 'RAMSI' stress-testing model to highlight these dynamics quantitatively. RAMSI uses disaggregated balance sheets covering the largest UK banks. For the simulations, we use data up to 2007 Q4 and draw 500 realisations from a macroeconomic model on a three-year forecast horizon to end-2010. The results highlight the role of contagion due to the systemic feedbacks. The distribution of total system assets at the end of the simulation period has a long left-hand tail, which is a direct consequence of the feedbacks, which can in some cases cause several institutions to default. This fat tail emerges in spite of the underlying shocks to macroeconomic variables having no such tail. These illustrative results point towards the importance of considering funding liquidity risk and systemic feedbacks in quantitative models of systemic risk.

The model could be extended in several ways. For example, rather than generating all shocks from a macroeconomic model, it would be interesting to allow for direct shocks to banks' cash-flow constraints, perhaps linked to some market-wide liquidity shock. It would also be helpful to capture the evolution of systemic liquidity crises incorporating more developed behavioural assumptions, and over a shorter time period than the three months used here. Finally, it would be interesting to use the framework to explore the role that macroprudential policies such as time-varying liquidity buffers might be able to play in containing systemic risk.

What do sticky and flexible prices tell us?

Summary of Working Paper No. 457 Stephen Millard and Tom O'Grady

Much recent research has looked at the microdata that make up price indices such as the UK consumer prices index (CPI). This work reaches three key conclusions. First, the microdata do support the underlying premise of the New Keynesian project, namely that there is a substantial amount of price stickiness. But second, underlying the headline inflation measures — which appear to be smooth and relatively autocorrelated (that is, current inflation is correlated with its own lags) — are inflation rates at the sub-component level that are much more volatile, and differ in terms of persistence. Third, and most importantly, the degree of price stickiness varies substantially across sectors. These results could potentially help us think about how inflation persistence arises. Inflation persistence may occur because the prices of different components of the CPI basket change at different speeds; some firms react to a shock immediately, whereas others take time to respond.

If that is the case, then prices that change at different speeds may also give us differing signals about the state of the economy. For example, relatively flexible prices may react more to the output gap than stickier prices: prices that change very frequently may be set on the basis of the current state of the economy. In contrast, relatively stickier prices may be more forward looking. If a firm knows that its price will last for a long time, it may think more about the future state of the economy when setting it. One implication is that sticky prices could tell us about firms' inflation expectations. Another is that we might want to look to flexible prices to see the impact of the output gap on inflation. And finally, the sticky component of inflation might be more useful than the aggregate for forecasting medium-run inflation, given that it drives persistence. This paper assesses these three claims against empirical evidence, and looks at how they hold up in the context of a formal model.

The paper first presents some empirical evidence that relatively flexible prices react more to deviations of output from trend than stickier prices, suggesting that prices that change very frequently are set on the basis of the current state of the economy. Some further evidence suggests that sticky prices contain information about firms' inflation expectations and that sticky price inflation may be useful in forecasting aggregate inflation two years out. These empirical results are then investigated further in the context of a dynamic stochastic general equilibrium model (which takes into account interactions between forward-looking optimising agents' choices in an economy subject to random shocks) containing a sticky price sector and a flexible price sector. Results generated by this model suggest that you would expect flexible price inflation to be more strongly related to the current output gap and sticky price inflation to medium-term inflation and inflation expectations, given standard economic theory.

Taken together, the results of this paper suggest that calculations of 'flexible price' inflation could, potentially, be used to provide monetary policy makers with a steer on the current state of the economy, in particular, the current output gap, which is notoriously hard to measure. In addition, calculations of 'sticky price' inflation could, potentially, be used to provide monetary policy makers with a steer on the medium-term inflation expectations of price-setters within the economy, again something about which it is hard to obtain any direct evidence.

A network model of financial system resilience

Summary of Working Paper No. 458 Kartik Anand, Prasanna Gai, Sujit Kapadia, Simon Brennan and Matthew Willison

The complex and opaque nature of modern financial systems poses a considerable challenge for the analysis of systemic resilience. An intricate web of claims and obligations links households and firms to a wide variety of financial institutions such as banks, insurance companies, and investment firms. The rapid development of securitisation and credit derivative markets has also made exposures between agents more difficult to assess and monitor in the absence of trade repositories. The global financial crisis illustrates how intertwined the financial network has become, while also making clear the potential for widespread losses and instability.

Recent efforts by central banks to measure and assess systemic risk have emphasised the important role played by network effects, fire-sale externalities, and funding liquidity risk in financial stability. A general insight is that these factors generate 'fat tails' in the distribution of aggregate losses for the banking system. That is, the financial system may incur very large losses with small probabilities.

Central bank studies typically rely on highly detailed, and relatively static, balance sheet data to establish precise linkages between banks in the domestic financial system and to derive banking system losses. This can be constraining when true linkages are not known (such as with credit risk transfer or off balance sheet activity) or when shocks strike financial institutions external to the core banking system. The pre-defined balance sheet interlinkages in these models also preclude analysis of how network structure matters for system resilience. The crisis has emphasised how network linkages and interactions between financial institutions are critical to understanding systemic risk. And the growing importance of 'stress-testing' exercises in the policy debate about financial stability points to the need for analyses that help overcome such limitations.

In this paper, we set out a general framework to gauge systemic risk in circumstances when data about the reach of financial exposures are limited and shocks are international in nature. We present a statistical model of a financial system involving a diverse set of financial agents, namely domestic banks, overseas banks, and firms, which are linked together by their claims on each other. We calibrate the model to advanced country banking sector data to illustrate how macroeconomic fluctuations, asset market liquidity and network structure interact to determine aggregate credit losses and contagion. Although the calibration is deliberately broad brush so as to emphasise the qualitative nature of the results, we obtain plausible loss distributions and can quantify, within the context of our model, the size of the macroeconomic or financial sector shock that may be necessary for system-wide failure to occur.

The model highlights how shocks are propagated through the direct interlinkages of claims and obligations among (and between) domestic banks and overseas banks. But it also shows how defaults across the network are amplified by asset fire sales and curtailed lending in the macroeconomy as 'credit crunch' effects take hold in the event of distress. In addition, we illustrate how greater heterogeneity of bank balance sheets leads to more realistic outcomes, characterised by the failure of some — but not all — banks in extreme scenarios. We also demonstrate how the model can be used to 'stress test' the banking system. The results obtained are entirely illustrative and only intended to demonstrate the usefulness of the framework.

Inflation and output in New Keynesian models with a transient interest rate peg

Summary of Working Paper No. 459 Charles T Carlstrom, Timothy S Fuerst and Matthias Paustian

Monetary models often assume that firms adjust their prices only slowly in response to changes in the economy. The most common assumption is that only some fraction of firms update their price in any period, as in the work of Calvo. At the same time firms and households are forward looking. They base their decisions not only on today's economic conditions but also on the outlook for the future, including expectations of future interest rates. The recent experience of low interest rates for an extended period of time provides a natural benchmark for testing such monetary models.

This paper examines the effects of an unconditional lowering of the nominal interest rate for an extended period of time in a model with infrequent price adjustment. One would expect that keeping the nominal interest rate low stimulates demand and thus increases inflation and output. In contrast, we show that the commonly used model of Calvo pricing implies unusual behaviour of inflation and output in such an environment. First, the anticipated unconditional lowering of the interest rate for an extended period of time can make initial inflation and output response from the model unusually large. Second, as interest rates are kept low for a sufficiently long period of time, inflation and output may actually fall in the model we consider. We show that this counterintuitive result is not simply due to using a linear model, but also obtains in the full non-linear sticky price model.

This is not an econometric test of sticky price models. Nor is it a statement about other possible shocks hitting the system. It is instead a question of *prima facie* plausibility. Our results suggest that these models can produce implausible behaviour under transient interest rate pegs. Future research should therefore examine whether similar results obtain under alternative models of price-setting or expectations formation.

Too big to fail: some empirical evidence on the causes and consequences of public banking interventions in the United Kingdom

Summary of Working Paper No. 460 Andrew K Rose and Tomasz Wieladek

Beginning in late 2007, the public sector around the world helped their struggling financial sectors in a number of different ways. Some banks were offered government funding or central bank liquidity insurance schemes, others received capital injections or were nationalised outright, and some were offered no support at all. To maintain future financial stability, it is important to not only understand the vulnerabilities that led the public sector to assist banks during the global financial crisis of 2008, but also assess the effectiveness of public sector help in stabilising individual banks' funding.

In the first part of this study, we therefore ask empirically what determined the style and recipients of public interventions. We use a confidential Bank of England bank-level data set using information on the balance sheets of all UK-resident banks. Our results suggest that the size of a bank is an important determinant of key public British banking interventions: capital injections, nationalisations, and government funding or central bank liquidity insurance schemes. In particular, the size of a bank relative to that of the entire banking system increases the probability of an intervention, suggesting that large banks are more likely to receive public sector assistance. This finding is consistent with the idea that some banks in the British banking system were deemed to be 'too big to fail'.

In the second part of this study, we study the consequences of public sector interventions in the British banking system. We argue that during the global financial crisis, financial institutions were subject to a bank run in wholesale markets. To improve our understanding of the effectiveness of these various public sector interventions, we study their effect on individual banks' wholesale to total liabilities ratio. Typically it would be difficult to credibly isolate cause and effect in our question of interest, since the banks that received government help were also the ones that were obviously most affected by a run on their wholesale liabilities. Fortunately, we established that bank size is an important determinant of government intervention in the first part of our investigation. This is a structural feature and changes only slowly over time. It is unlikely to be affected by sudden movements in bank liabilities and can be used to predict government intervention. We therefore use a bank's relative size with respect to the whole banking system to isolate the causal effect of British public sector interventions on an individual bank's wholesale funding. We find that these interventions mattered in a tangible sense: they seemed to restore access to wholesale funding. More precisely, the share of wholesale (non-core) funding rose significantly following intervention. As one objective of UK public sector intervention during the global financial crisis was precisely to stabilise flighty financial market funding, it seems to have been effective.

Labour market institutions and unemployment volatility: evidence from OECD countries

Summary of Working Paper No. 461 Renato Faccini and Chiara Rosazza Bondibene

The interest in how unemployment responds to output fluctuation is long-standing. Okun's rule of thumb, the empirically observed relationship between changes in unemployment and changes in output, has been a useful guide for policymakers since it was proposed in 1962. However, the relationship between unemployment and output is not stable over time and differs markedly across countries. Despite its importance, the factors underlying the response of unemployment to output fluctuations are not well understood. We investigate whether laws regulating the labour market, typically referred to as labour market institutions, can help explain cross-country and time variation in this relationship.

The sensitivity of unemployment to cyclical changes in output differs considerably across OECD countries and has changed over time in most cases. In particular, the United States, together with other Anglo-Saxon economies including the United Kingdom, Canada and New Zealand, are characterised by average values of the sensitivity of unemployment to output movements. Compared to the Anglo-Saxon economies, unemployment volatility is lower in Mediterranean countries and higher in Scandinavian economies. So can differences in labour market institutions account for this heterogeneity?

We find that labour market institutions are indeed an important factor affecting the response of unemployment rates to changes in output. The impact of most labour market institutions is found to be statistically significant; for some, such as employment protection and unions, the quantitative impact is particularly strong. In particular, we find that employment protection strongly reduces the cyclical response of unemployment. In addition, we find that the precise nature of union bargaining has important consequences for cyclical unemployment dynamics: while union coverage (the proportion of employees covered by collective agreements) significantly increases fluctuations of unemployment rates, union density (the proportion of employees who are members of a trade union) has the opposite effect. Our interpretation of these findings is that union bargaining generates real wage rigidities, whose impact increases with the spread of union agreements (union coverage). As a consequence of stronger real wage rigidities unemployment becomes more volatile. However, unions also care about job security for their members and therefore the sensitivity of unemployment fluctuations to changes in aggregate production will decrease with density. These findings are consistent with previous results showing that union coverage is positively related with downward real wage rigidities and with evidence that union membership decreases the probability of dismissals.

The benefit replacement rate, the duration of unemployment benefits and taxation are found to have a limited impact on the sensitivity of unemployment fluctuations. All of these institutions appear to slightly reduce the cyclical response of unemployment. These results do not support the theoretical predictions that labour market institutions could increase the volatility of unemployment by reducing the available 'surplus' divided by firms and workers in the bargaining process.

Overall, we find that institutions explain about one quarter of the explained variation, which in turn is about half of the total observed variation. So, we conclude that labour market institutions are an important factor governing cyclical unemployment fluctuations, but they are not the entire story. Finally, we find some evidence supporting the hypothesis that interactions between shocks and institutions matter for cyclical unemployment dynamics.