

# Exploiting the monthly data flow in structural forecasting

**Summary of Working Paper No. 509** Domenico Giannone, Francesca Monti and Lucrezia Reichlin

---

Central bankers typically place their forecasts centre-stage in the communications. The forecast is presented as a guide to future developments in the economy, but also as a means of communicating policy. Thus statistical models that might forecast well but have no economic interpretation are insufficient to meet the policy explanation task. Yet to deliver accurate predictions, especially in the short term, it is vital that the forecast incorporates a range — often a very wide range — of timely information.

Consequently the preparation of a forecast is a subtle process of aggregation of knowledge of how the economy and policy channels work and detailed data. This process involves the combination of formal models which aid understanding, expert judgement and statistical data analysis. In this paper we address a particular part of this process and analyse the connection between two important tools in the forecasting process: the structural quarterly model and the daily monitoring of monthly data releases for the assessment of the current state of the economy.

The structural quarterly model is often a dynamic stochastic general equilibrium (DSGE) model which looks at the joint evolution of the key relations in the economy as it is buffeted by random (stochastic) shocks and is deeply driven by economic theory. It is essential when constructing scenarios based on different policy paths or other conditioning assumptions; that is, for policy analysis. The objective of this is not to obtain a simple forecast, but rather to analyse the implications of policy alternatives. Moreover, from structural models one can recover quantities that are not directly observable from the data but that are often relevant for the understanding of the stance of policy, such as the natural rate of interest or potential output. Although this part of the analysis is essential for guiding the policy discussion, any decision-maker needs to have, in addition, a system in place for understanding the evolution of the current state of the economy. Such a system involves the analysis of many different items of data, including surveys or conjunctural leading indicators which are published early in the quarter, before the release of the quarterly national accounts data that

the DSGE model is largely designed to explain, and can provide a timely signal on quantities of key interest such as GDP or employment. For this function, the typical structural model is of no use since it is not designed to capture realistic features of the data flow: a non-synchronous calendar of publications, mixed frequency (meaning a mixture of quarterly, monthly or even daily data), and potentially a very large dimension, with perhaps hundreds of data series. But recent work has developed a statistical framework for dealing with these problems, allowing continuous updates of the estimate of the current state of the economy in relation to the real-time data flow. This process is sometimes labelled nowcasting, the point being it is helpful in very short-run forecasts.

This paper proposes a framework that bridges a structural quarterly model and a statistical model for nowcasting. This is particularly relevant for the conduct of monetary policy today when, with the implementation of forward guidance, an increasing emphasis has been placed on the definition and communication of nearer-term policy in relation to the evolution of the state of the economy. With our methodology it is possible to exploit the real data flow in a DSGE model, using the most recent data for both the variables that appear in the model and other series that are not explicitly modelled, but that might be informative about the current state of the economy. Because this information is processed through the structural model, we can update the policy analysis at each new data release and assess the impact of the new piece of information on our scenarios and on our estimates of policy-relevant and inherently model-dependent concepts, such as the equilibrium real interest rate or potential output.

We apply our methodology to a medium-scale DSGE model and to a sizable number of timely macroeconomics series for the United States and we show that we obtain point and density nowcasts that are much more accurate than those obtained with the quarterly model, and comparable with a traditionally tough benchmark, such as the Survey of Professional Forecasters' nowcasts.

# Institutional investor portfolio allocation, quantitative easing and the global financial crisis

## Summary of Working Paper No. 510 Michael A S Joyce, Zhuoshi Liu and Ian Tonks

In its efforts to loosen monetary conditions in March 2009 in response to the effects of the deepening financial crisis on the inflation outlook, the Bank of England's Monetary Policy Committee (MPC) reduced policy rates to their effective lower bound of 0.5% and began a programme of large-scale purchases of financial assets financed through the creation of central bank reserves, so-called quantitative easing (QE).

While the objectives of the QE policy were clear, there has been more debate over how the policy was expected to work. The MPC has often emphasised the portfolio balance channel as a key element in the transmission of the Bank's asset purchases to the rest of the economy during the financial crisis. According to this mechanism, purchases of financial assets from the non-bank private sector (eg insurance companies and pension funds) financed by central bank money initially increase broad money holdings and push up asset prices, as those who have sold assets to the central bank rebalance their portfolios into riskier assets. This then stimulates expenditure by increasing wealth and lowering borrowing costs for households and companies.

Despite the emphasis put on this channel by monetary policy makers, the role of portfolio balance effects in theory remains rather controversial and the empirical evidence in its support rather indirect. Most empirical research on the topic has inferred the importance of this channel from the behaviour of government bond prices/yields and other asset prices, rather than from direct evidence on the behaviour of investors. Although discussions with market contacts have also suggested some asset reallocation occurred in response to QE, particularly during the first phase of asset purchases between March 2009 and January 2010, there has been little hard evidence on the behaviour of insurance companies and pension funds, the group of investors who the Bank specifically targeted by mainly purchasing medium to long-term UK government bonds (one of the main assets held by these institutions).

The aim of this work is to try to fill this gap. We examine the behaviour of institutional investors, ie insurance companies (particularly life companies) and pension funds, both before and during the crisis and whether their portfolio allocation behaviour is consistent with portfolio balance effects. If QE has partly worked through a portfolio balance channel then we would expect that institutional investors will have reduced their holdings of UK government bonds (gilts) below what they would otherwise have been and that they will have increased their demand for riskier assets. This of course raises the difficult issue of inferring what would have happened in the absence of QE (the 'counterfactual').

In order to generate a plausible counterfactual, it is clearly important to allow for a range of other factors that may have been relevant in driving portfolio allocation. At the same time, allowing for the influence of other factors that may have been influenced by QE (eg domestic financial conditions) may lead to understating the potential effects of the policy (a switch into riskier assets will be attributed to improved financial conditions rather than QE, even though the policy may have been behind the improvement). We address this issue by allowing only for factors that influence portfolio allocation, but at the same time are unaffected by the Bank's purchases. These factors include gilt issuance by the Debt Management Office and foreign financial variables. To measure the impact of QE, we use in and out-of-sample model-based forecasts to construct counterfactuals of what would have happened if the policy had not been implemented.

Our analysis of a range of data sources, including national accounts net investment data and micro-level data on individual life insurance companies and pension funds, is consistent with QE having led institutional investors to shift their portfolios away from gilts towards corporate bonds relative to the counterfactual. Analysis of the microdata shows that the switch into corporate bonds apparent in the sectoral data is reflected in remarkably similar behaviour across different types of life insurance companies and pension schemes, but in the case of insurers (who appear to exhibit more heterogeneity) the switch away from gilts was more pronounced for companies that showed less risk aversion (ie were larger than average and more heavily weighted in equities), were under more financial constraints (ie had a lower-than-average ratio of business premiums to assets) and those less constrained on average by their liabilities (ie with a larger share of assets linked directly to liabilities). For pension funds, the switch out of gilts was more pronounced for those funds that were better funded.

Overall the balance of our evidence is consistent with the hypothesis that the Bank of England's QE policy resulted in some portfolio rebalancing behaviour by institutional investors, who appear to have reduced their gilt holdings and reinvested some of the proceeds into corporate bonds relative to the counterfactual. But it appears that portfolio rebalancing was limited to corporate bonds, with most of the evidence suggesting that institutional investors moved out of equities during the period of QE purchases. Of course, this does not necessarily imply equity prices were not supported by portfolio reallocation behaviour, still less from QE, as our analysis only considers insurers and pension funds and we do not investigate the behaviour of other financial institutions; something we leave for further work.

# QE and the bank lending channel in the United Kingdom

## Summary of Working Paper No. 511 Nick Butt, Rohan Churm, Michael McMahon, Arpad Morotz and Jochen Schanz

In March 2009, the Bank of England's Monetary Policy Committee (MPC) voted to commence a programme of asset (predominantly gilt) purchases, commonly referred to as quantitative easing (QE). Following subsequent rounds of purchases the stock of asset purchases reached £375 billion by October 2012. Faced with a likely deep recession and the risk of deflation, this policy was intended to boost GDP and inflation. The MPC thought that this would primarily happen by QE reducing gilt yields and boosting the price of a range of assets. That view drew on the monetary economics literature, which suggests that when sellers of gilts — who were primarily other financial corporations (OFCs) such as pension funds, insurance companies and asset managers — receive deposits, they would wish to rebalance their portfolios in to riskier assets, due to money and securities not being close substitutes. There is a broad range of evidence that suggests that QE did reduce gilt yields and boost other asset prices. The economics literature also suggests that such expansionary monetary policy may lead to a shift in banks' willingness to lend, via a 'bank lending channel'. At the time QE was launched the MPC were not expecting or relying upon a large bank lending channel due, in part, to the pressures on banks to decrease the size of their balance sheets. In this paper we test whether QE did in fact provide a boost to bank lending.

We show, using a simple framework, that changes in banks' liabilities resulting from QE are likely to lie at the heart of any bank lending channel. The key insight is that a shock that boosts banks' OFC deposit funding can lead to a greater willingness of banks to lend, as these deposits offer a cheaper source of financing than other sources of funding. But if the variability, or 'flightiness', of these deposits increases then banks are less likely to increase their lending at a given price, as cheaper funding today may have to be replaced with more expensive funding tomorrow.

We use this framework to inform our empirical analysis, which makes use of a data set available to researchers at the Bank of England. It combines balance sheet, regulatory and market operations data for individual banking groups. This allows a descriptive review of banks' balance sheets over the QE period. We find that banks that took part in gilt sales saw increases in reserves and OFC deposit positions but that only a portion of the proceeds remained at the end of the month. Indeed, we also show that the variability of banks' deposit and reserve positions increased during QE, which could be consistent with the portfolio rebalancing channel of QE. These findings help inform our empirical tests of the bank lending channel and our interpretation of them.

A key challenge for empirical work on the bank lending channel is to isolate changes in lending caused by changes in deposits, from changes in deposits caused by new lending (an endogenous variation in deposits). We attempt to address this problem using two alternative approaches. Our first approach exploits the fact that, for historical and infrastructural reasons, it is likely that not all banks are equally well placed to receive very large OFC deposits. We use historical data on the share of banks' OFC funding (relative to their balance sheet) to identify a group of banks that are most likely to have received deposits created by QE, which we call 'OFC funders'. We use this variable, along with variation in banks' OFC deposit funding to test whether there was a bank lending channel by comparing the lending response of such OFC funders to that of other banks during the QE period. Our second approach makes use of the fact that while most gilt purchases were from OFCs, these had to be settled via banks who were market makers in gilts. As these gilt sales were likely to be unrelated to banks' lending decisions, we can use data on gilt sales to remove the endogenous variation in banks' OFC deposit holdings and so test for a bank lending channel using an instrumental variables approach that controls for the interrelatedness of the bank's decision.

We find no statistically significant evidence from either approach that those banks who received increased deposits from QE lent more, all else equal. This contrasts with results in the recent Bank of England working paper, 'Quantitative easing and bank lending: a panel data approach', which used different empirical methods placing more weight on the pre-crisis relationship with deposits to show that QE may have had a statistically significant but small effect on bank lending. Our results do not preclude a bank lending channel, but if the effect were very powerful it seems unlikely there would be no evidence of it in our tests. While our results do not provide an explanation of why such a channel did not operate, our framework suggests that if QE gave rise to flighty deposits, then the traditional bank lending channel would be diminished. And our analysis suggests that QE has been associated with an increase in the variance of banks' reserves and OFC deposit positions. This is consistent with the idea that there was no bank lending channel from QE precisely because portfolio rebalancing was occurring and is therefore consistent with other studies which show that QE boosted aggregate demand and inflation. UK policymakers did not rely on QE to boost bank lending and our evidence lends support to the use of other policies, rather than QE, to attempt to improve the supply of credit.

# Policy uncertainty spillovers to emerging markets — evidence from capital flows

**Summary of Working Paper No. 512** Ludovic Gauvin, Cameron McLoughlin and Dennis Reinhardt

Since the end of the 'Great Moderation' and the global financial crisis, policymakers have discussed the appropriate policy mix for returning to sustainable growth. A marked feature of this discussion has been the effects of macroeconomic policy uncertainty on domestic investment decisions by firms, especially in the light of the uncertain US fiscal outlook and the ongoing euro-area crisis. At the same time, concerns regarding the impact of domestic policies on other economies — ie 'spillover effects' — feature prominently in the international policy debate. In particular, attention has focused on the spillover impacts of capital control policies, as well as the external impact of monetary policy settings in advanced countries. More recently, these two debates have been drawn together to analyse the spillover effects of advanced-country policy uncertainty to investment and output in the rest of the world.

In this paper, we examine whether such policy uncertainty spillovers have been transmitted via cross-border capital flows. Specifically, we examine whether macroeconomic policy uncertainty in the United States or the European Union (EU) spilled over to emerging market economies (EMEs) via gross portfolio equity or bond inflows.

In principle, policy uncertainty could lead to an increase or decrease in portfolio inflows to EMEs. On the one hand, a less predictable political environment would tend to hinder domestic growth prospects, decreasing the attractiveness of investing in a given country (recent evidence points to effects of policy uncertainty on domestic output and investment). Based on this we would — *ceteris paribus* — expect investors' to shift more of their investment abroad given the declined attractiveness of investing in the United States or the EU. On the other hand, higher policy uncertainty may decrease the overall size of investors' positions in relatively more risky investment funds. Since there is a strong relationship between macroeconomic policy uncertainty and the US equity risk premium, higher policy uncertainty may impact advanced-economy investors' willingness to take risk and lead to safe-haven flows (consistent with outflows out of EMEs that are often perceived as less 'safe'). In response to an investor funding shock, funds considerably change their allocations to emerging markets. Our paper can be interpreted as assessing the relative strength of these competing hypotheses for policy uncertainty shocks originating from two distinct regions and distinguishing between bond and equity inflows.

We find — using first a linear regression framework — that increases in policy uncertainty in the United States tend to significantly reduce both bond and equity inflows into EMEs. Conversely, increases in EU policy uncertainty tend to have different effects on equity versus bond flows into EMEs: bond inflows into EMEs decrease, but equity flows to EMEs increase in response to increased EU policy uncertainty. This is consistent with the hypothesis that shocks to US policy uncertainty are associated with safe-haven equity flows out of EMEs whereas the reduced attractiveness of investing in the EU following shocks to

EU policy uncertainty appear to outweigh any safe-haven equity flows out of EMEs.

Non-linearities play, however, an important role in the size and direction of spillover effects. First, we provide evidence for two structural breaks in the relationship between changes in policy uncertainty and capital flows. The first break coincides with the first large increases in the cost of insuring against mortgages of lower credit ratings (BBB- and BBB) in the United States, providing evidence that the impact of the coming financial crisis was felt in portfolio flows slightly before the onset of funding illiquidity in the interbank market. The second break occurs in November/December 2010, coinciding with a significant expansion of QE2 by the US Federal Reserve in November 2010.

The level of global risk performs best in explaining non-linearities. Increases in EU policy uncertainty have a significantly negative impact on bond inflows only in the high global risk regime and, pointing into the same direction, the spillover impact of EU policy uncertainty on equity inflows is less positive in the high global risk regime than the low global risk regime. Global risk (proxied by the VIX index in our baseline) appears therefore not only as an important determinant of capital flows on its own, but it also determines how other push/pull factors (including policy uncertainty) impact portfolio flows. Turning to domestic factors, we find that the impact of policy uncertainty on bond inflows does not depend on domestic variables: changes to policy uncertainty have for example the same impact on bond inflows independent of a country's level of sovereign risk or equity market returns. Conversely, the level of country-specific sovereign risk (as proxied by credit default swap spreads) does determine the magnitude of policy uncertainty spillovers via equity flows. Increased EU policy uncertainty pushes portfolio equity inflows into EMEs even if global risk is high, but only into countries with low sovereign default risk.

Portfolio flows from funds based in the United States may show different reactions to EU policy uncertainty shocks than portfolio flows from funds based in the EU itself. The degree of home bias may play a crucial role. And to the extent that policy uncertainty with regard to macroeconomic policies impacts variables such as investors' wage income risk, it may also affect fund investors' willingness to buy risky assets, including assets held in EMEs. Accounting for the domicile of funds does, however, not change the finding on the positive spillover impact of EU policy uncertainty on equity flows to EMEs: we observe positive spillover effects in both the low and high global risk regime even for equity flows originating from funds domiciled in the European Union. These spillover effects are stronger for equity flows originating from funds domiciled in the United States: in the high global risk regime, flows into EMEs from US-domiciled funds increase even into EMEs with high sovereign default risk, whereas, mirroring our aggregate results, flows from Europe-domiciled funds increase only into EMEs with low sovereign default risk.

# Variations in liquidity provision in real-time payment systems

## Summary of Working Paper No. 513 Edward Denbee, Rodney J Garratt and Peter Zimmerman

A payment system consists of the procedures and associated computer networks used by its participants to transfer money. Sometimes called the 'plumbing' of the financial system, smoothly functioning payment systems are essential to the operation of financial markets. Large-value wholesale payment systems, such as CHAPS in the United Kingdom, are generally considered to be systemically important because of the value and nature of the financial transactions that they facilitate. On a typical business day, transactions with a total value of around £277 billion flow through CHAPS, roughly equivalent to one sixth of the United Kingdom's annual gross domestic product.

In a real-time gross settlement (RTGS) system, payments settle immediately and with finality in central bank money, providing that the paying bank has sufficient liquidity to fund the outgoing payment. But the aggregate amount of liquidity needed to fund payment obligations is often much less than gross payment flows. That is because, during the course of the day, each bank in the payment system typically makes and receives thousands of payments. Thus outgoing payments are not only funded from liquidity made available from banks' own reserves, but also from liquidity obtained from incoming payments, which can be recycled to fund a bank's own outgoing payments.

If banks were required to process payment requests as soon as they received them, then they would have little discretion over the liquidity they provide to the rest of the payment system. But this is not usually the case: with a few exceptions, banks do not usually have to process payment requests as soon as they receive them. Rather, banks may choose to delay processing payments in order to conserve liquidity and to make use of money from recycled incoming payments. They may do this because using their own liquidity has an opportunity cost in terms of foregone investment opportunities, or to help mitigate against the risk of liquidity shocks later in the day. But if too many banks withhold liquidity the payment system can fall into gridlock, preventing payments from being made, with consequences for financial stability. This is why central banks have an interest in monitoring banks' liquidity provision in order to ensure the continued smooth functioning of the payment system.

In this paper, we measure liquidity provision in two ways. First we look at the maximum net debit position that banks in the payment system reached in their accounts over the course of each day, during a historic period. The sum of these net debit positions across all banks is the total amount of liquidity that was actually used to make the day's payments. Therefore each bank's own net debit position, divided by the sum of the net debit positions of all banks, gives the share of liquidity provided by each bank. Whenever the value of a bank's payments into the system exceeds that of those it has received, the difference has to be made up either from central bank reserves, or from eligible collateral that a bank pledges intraday in order to obtain liquidity from the central bank. This means we can assume that a net debit position imposes an opportunity cost of using central bank reserves or of pledging eligible collateral, and so our first measure attempts to measure the nominal monetary cost of liquidity provision.

Our second measure examines a bank's exposure to counterparty risk. Making payments earlier can result in a greater exposure to counterparty risk, in addition to the monetary cost described above. For example, if the paying bank relies on recycled liquidity to fund future payments, then it faces the risk that its counterparty fails to recycle the liquidity back into the payment system in a timely fashion. This may happen, for example, if the counterparty has an operational problem or enters bankruptcy. We capture this using a bank's average net debit position throughout the day.

We compute our measures for CHAPS member banks using data from January 2008 to May 2010. Since some banks have a higher value of payment activity than others — and hence may reasonably be expected to provide more liquidity in absolute terms — we adjust our measures using banks' shares of payment activity. We present our aggregated results for groups of banks in two size categories. Although larger banks do provide the bulk of the liquidity in absolute terms, we find that the smaller banks almost invariably provide a larger share of liquidity to the system than their share of payments. This is true under both of our measures.

We use a Gini coefficient measure to capture variations in the liquidity cost of payments among CHAPS banks. We find that the variations significantly increase in the period surrounding the collapse of Lehman Brothers. Unobserved factors — such as differences in the nature of the payments activity — may explain some of these differences in liquidity provision. In any case, some heterogeneity between individual banks' liquidity provision and usage is inevitable, and does not necessarily imply unfairness. Since payment requests from customers often arrive outside of the control of the banks, there will be net liquidity providers and users on any given day, even if all banks process payment requests immediately. This means the patterns of liquidity provision that we observe could simply reflect the way payment requests happened to arrive.

But we would like to know when observed differences in liquidity provision are so marked that they are very unlikely to have solely been a result of external factors. We provide a method for identifying when the observed variations in liquidity provision are unlikely to have occurred by accident. The idea is to ask, given all the different permutations for how payments might have arrived, what would be a very unlikely level of liquidity provision? We answer this question by reshuffling each day's payment schedule 200 times to generate distributions for our liquidity provision measures. We then check how often actual values of these measures lie in the tails of these distributions. We find that instances where banks are in these tails occur far more frequently than we would expect to see in the absence of behavioural or structural factors.

Our methodology can be applied to any RTGS system to investigate patterns of liquidity provision at the level of individual payment system members. In the particular case of CHAPS, it should be noted that the system has undergone several structural changes since the end of our data period in May 2010 that may have led to changes in the patterns of liquidity provision.

# Optimal monetary policy in the presence of human capital depreciation during unemployment

## Summary of Working Paper No. 514 Lien Laureys

---

The recession and associated rise in unemployment has helped to revive interest in studying the trade-off that monetary policy makers face between unemployment and inflation stabilisation. But the literature has focused primarily on an environment where all workers have the same characteristics, leaving it an open question whether this trade-off is altered once worker heterogeneity is taken into account.

This paper analyses this trade-off in an environment where human capital depreciation during unemployment generates heterogeneity among *ex-ante* identical workers. This source of heterogeneity seems particularly relevant because when workers are exposed to human capital depreciation during periods of unemployment, job creation affects the unemployment pool's composition in terms of skills, and hence the economy's production potential. If aggregate shocks induce changes in the skill composition of the unemployment pool which are not desirable from a social point of view, it might be optimal to influence job creation by allowing for more or less inflation relative to an environment where human capital depreciation is not taken into account. Put differently, the presence of skill erosion during unemployment might affect the trade-off between unemployment and inflation stabilisation.

In models where the unemployed search for jobs, known as matching models, there are two market failures that lead to inefficiency. An unemployed searcher takes into account the personal costs and benefits of search but ignores the effects on others, giving rise to a congestion externality. The more an unemployed worker searches the easier it is for firms to hire

which means more production on average. But at the same time it becomes harder for the other unemployed workers to find a job (there is a similar effect with vacancies). There is a point at which the costs cancel out — the Hosios condition — so in this case an economy where there are flexible prices is nevertheless optimal.

Introducing human capital depreciation during unemployment into an otherwise standard New Keynesian model incorporating sticky prices and other features including search frictions in the labour market leads to the finding that the flexible-price allocation is no longer efficient even when the Hosios condition holds. This is because it generates an additional composition externality in job creation: firms ignore how their hiring decisions affect the extent to which the unemployed workers' skills erode, and hence the output that can be produced by new matches. Consequently, it might be optimal (meaning welfare maximising) for monetary policy to deviate from strict inflation targeting (which in this simplified model means that the policymaker will always try to hit the inflation target in every period, thus mimicking the flexible price equilibrium).

In the paper a theoretical model incorporating this mechanism is calibrated using standard values so that it is broadly consistent with the benchmark US data. It emerges that optimal price inflation is no longer zero. But deviations from it are almost negligible. Consequently, the prescription for the conduct of monetary policy does not change much when it is taken into account that the unemployed are exposed to human capital depreciation: optimal monetary policy stays close to strict inflation targeting.

# The Bank of England *Credit Conditions Survey*

## Summary of Working Paper No. 515 Venetia Bell and Alice Pugh

Credit conditions play an important role in the transmission of monetary policy. The tightening in credit conditions during 2008 and 2009, when lending growth fell and lending rates rose, reduced the impact of the loosening in monetary policy during this period. More recently, credit conditions have relaxed as the recovery has started to take hold.

The Bank of England's *Credit Conditions Survey (CCS)* is a survey of UK lenders, introduced in 2007 in order to provide a better understanding of developments in credit markets. The CCS has subsequently been used heavily by policymakers and economic commentators. It had three main aims:

- To provide a better understanding of developments in credit markets. As an example, it could be used to help explain whether the weakness in bank lending observed after 2008 reflected a tightening in the supply of credit, or a weakening in the demand for credit.
- To collect information on all the components of the cost of a loan. In particular, the CCS asks lenders for information on fees, non-price terms and quantity restrictions, each of which may be important in determining loan volumes.
- To provide an early steer on developments in credit conditions, by asking about lenders' expectations of developments three months ahead.

This is the first paper to provide a detailed econometric analysis of individual banks' responses to the CCS. We investigate two main questions: how well do the responses track movements in the official rates and lending data, and are they useful for predicting changes in credit spreads and lending growth one quarter ahead?

The first of these questions is important because some of the information on credit conditions collected in the CCS is otherwise unobserved. If the survey responses correspond closely to movements in the official quantitative data, where comparable data exist, this would give us some confidence that the survey responses are also informative about these unobservable movements in credit conditions. The second question should help us to determine whether or not banks'

survey expectations contain additional information (over and above existing data) which is useful for predicting changes in credit conditions in the near term.

We combine individual banks' responses from the CCS with bank-level microdata on lending growth and credit spreads. The use of individual bank data rather than aggregate data allows for an increased number of observations on which to perform the estimation, and allows us to relate changes in credit conditions to the same panel of banks over time.

We find that the survey responses match available official data from other sources. Over the period 2007 Q2–2013 Q2, many of the responses are significantly associated with changes in credit spreads and lending growth. But results vary by type of lending. The relationship between the responses and official data is strongest for household lending, and for headline or popular loan products. Responses relating to corporate lending are less well correlated with the official data, although this may be due in part to a lack of suitable official data with which to compare the responses.

In the second part of the analysis, we find that the CCS contains some predictive power for credit spreads and lending growth one quarter ahead, although results vary by type of lending. Changes in spreads on two-year fixed-rate mortgages can be partially explained by lenders' survey expectations of changes in spreads and default rates over the subsequent quarter. And survey expectations of looser credit availability and credit scoring criteria have some predictive power for lending growth in the subsequent quarter.

Overall, the results contained in this paper suggest that the CCS contains useful empirical information for policymakers with respect to developments in credit conditions. And while we have only investigated the relationship between the responses and quantitative data where comparable quantitative data exists, the results provide grounds for believing that the CCS gives a useful steer for aspects of credit conditions that are not otherwise observed. Nevertheless, the short sample period means that the results are necessarily preliminary and should therefore be treated with some caution.

# Mapping the UK interbank system

## Summary of Working Paper No. 516 Sam Langfield, Zijun Liu and Tomohiro Ota

---

This paper maps the structure of the network of interbank connectedness in the UK banking system. Using a new regulatory data set on the UK interbank exposures, we construct two networks: the exposures network is comprised of banks' counterparty credit exposures to other banks across different financial instruments; and the funding network aggregates banks' cash funding from other banks.

The exposures network and the funding network have different structures. The exposures network exhibits a 'core-periphery structure', in which core banks are densely connected to each other and peripheral banks are weakly connected to each other. The derivatives market in particular is characterised by a densely connected core, which we interpret as evidence of there being strong economies of scale associated with trading derivatives. In contrast, the funding network has less of a core-periphery structure, owing to a lower degree of connectedness among core banks in the unsecured lending and repo markets.

These structural differences between the two networks suggest that credit risk and liquidity risk propagate in the interbank system in different ways. To dig deeper, we divide

banks into clusters according to the markets in which their interbank activity is concentrated. Large derivative houses dominate the system, absorbing funding from all other clusters, particularly non-UK investment banks (using repo) and smaller UK banks (using unsecured loans). A reduction in funding provided by these banks could trigger widespread liquidity shortages.

We also identify contagious links, where a bank's single counterparty exposure is greater than its capital. We identified the contagious links from core banks to many peripheral banks, implying that the isolated default of certain core banks causes multiple peripheral banks to default. However, higher-round effects from these defaults appear to be relatively limited, given that core banks tend to be relatively well diversified with respect to their bank-counterparty credit risk. We infer that core-periphery structures tend to be robust, because core banks can act as fire-stops against contagion. But such structures are also potentially fragile, because a core bank's distress could propagate throughout the network. In principle, this finding supports the application of capital surcharges on systemically important banks to build the resilience of these fire-stops in the core of the network.



# Optimal contracts, aggregate risk and the financial accelerator

## Summary of Working Paper No. 517 Timothy S Fuerst, Charles T Carlstrom and Matthias Paustian

---

Frictions in credit markets are widely known to amplify business cycles. The mechanism typically works via leverage (the ratio of debt to net worth) of the borrower. Typically, an adverse macroeconomic shock reduces the value of the assets of credit-constrained borrowers. The resulting fall in borrowers' net worth increases leverage. In turn, higher leverage makes an underlying credit friction more severe and raises credit spreads. As a consequence, demand for investment falls by more than would happen in a world without credit market frictions, depressing asset values further. This sets in motion a feedback loop between rising spreads and falling asset prices that is at the heart of the financial accelerator.

The most prominent paper incorporating the financial accelerator mechanism in a quantitative macroeconomic model was published in 1999 by Ben Bernanke, Mark Gertler and Simon Gilchrist (BGG). We revisit the debt contract they employed and highlight how the financial accelerator depends on the treatment of aggregate risk in the debt contract. BGG study the optimal financial contract in a world where borrower and lender have asymmetric information about firm-specific productivity. Lenders can only observe the return of the firm's project by paying a monitoring cost. In addition, there is aggregate macroeconomic risk that is costlessly observable by everyone. The key assumption is that the return to the lender does not depend on the realisation of aggregate risk.

In this paper, we derive the optimal financial contract and show that the return to the lender varies with the realisation of aggregate risk. Consequently, the interest rate in the optimal debt contract is contingent on aggregate macro

variables, much as the coupon payment in an inflation-indexed bond is linked to the particular realisation of aggregate inflation. This 'state contingency' in the optimal contract is, however, rather complex. We show that the lender return varies with shocks to household consumption, the aggregate return on capital and the marginal value of internal funds of the borrower.

A key feature of the state-contingent debt contract is that it limits fluctuations in leverage and greatly reduces the financial accelerator. When an unexpected adverse macroeconomic shock reduces the return on borrowers' investments, the loan contract calls for a reduction in the borrowers' interest rate. As a result, fluctuations in net worth and leverage are limited and much of the adverse feedback loop described above is avoided. Ultimately, aggregate risk is shared between households (lenders) and entrepreneurs (borrowers), rather than falling predominantly on the borrowing-constrained firm as in BGG. In a model calibrated to match US data, we show that this contract implies a welfare improvement for both parties. Furthermore, amplification from credit frictions is negligible.

It is an open question to what extent actual contracts are state contingent in the way our analysis suggests. At face value, it seems that such contingency is very rare. Our primary contribution is to derive the optimal debt contract in the BGG model, not to state that financial frictions in the data cannot amplify macro shocks. But the analysis also enables us to quantify the welfare cost of financial frictions. We find that the costs of frictions are small, increasing in adjustment costs.