Shadow banks and macroeconomic instability

Summary of Working Paper No. 487 Roland Meeks, Benjamin D Nelson and Piergiorgio Alessandri

Between the early 1990s and the onset of the 2007–09 sub-prime crisis, the financial system in the United States and elsewhere underwent a remarkable period of growth and evolution. Banking underwent a shift away from the traditional 'commercial' activities of loan origination and deposit issuing towards a 'securitised banking' business model, in which loans were distributed to entities that came to be known as 'shadow' banks. As shadow banks came to replicate core functions of the traditional banking system, in particular those of credit and maturity transformation, they took on many of the same risks but with far less capital. An overreliance on securitisation, and the increased leverage of the financial system as a whole, ultimately contributed to financial instability, recession, and a substantial contraction in shadow banking activity.

The aggravating role played by flaws in the securitised banking model have been rightly emphasised in many accounts of the sub-prime crisis and ensuing 'Great Recession'. But there is also a need to understand the increasingly central role played by securitisation in credit provision over the decades prior to the crisis. To illustrate why, we show that in the US data from 1984 to 2011 periods when traditional bank credit underwent cyclical contraction were often periods when shadow bank credit expanded. Similarly, other authors have documented that over the post-1984 period, consumer credit and mortgage assets held by commercial banks were positively correlated with GDP, while holdings outside the banking system were negatively correlated with GDP.

These observations suggest that a macroeconomic model which seeks to account for the behaviour of intermediated credit should be able to account for the differences in credit supply across institutions, as well as the collapse in shadow banking during the crisis. To that end we develop a dynamic general equilibrium model featuring securitisation and shadow banking, which aside from its treatment of the financial sector, closely resembles a standard macroeconomic model.

In our model we show that the ability of commercial banks to securitise can stabilise the overall supply of credit in the face of aggregate disturbances, but that risk-taking by the shadow banking system leads to an increase in macroeconomic volatility. We then give conditions under which the negative correlation between traditional and shadow bank credit observed in the US data come about, and quantify the credit dynamics resulting from the interaction between banks and shadow banks. Finally, we argue that in a securitisation crisis government policies targeted at the shadow banking system, such as purchases of asset-backed securities, can have spillover effects on the rest of the financial system which weaken the effectiveness of interventions. Taken together, these points constitute a first step towards addressing what are widely thought to be some important shortcomings of the generation of dynamic general equilibrium models used for research and policy analysis prior to the recent crisis.

Our model does not attempt to capture the full complexity of shadow banking activities, however, and leaves room for future research. First, we do not attempt to model the process of financial innovation and regulatory change which lay behind the rapid expansion of shadow banking. Second, the crisis highlighted shortcomings in the workings of key asset markets, which we ignore. For example, we do not model complex financial instruments based on securitised assets, such as collateralised debt obligations, which the market badly mispriced. Last, we do not deal with issues of prudential regulation, or with policies relating to financial system structure. An important contributory factor behind the creation of some shadow banking entities, in particular structured investment vehicles, was a desire by banks to reduce the amount of regulatory capital they held against credit exposures. In our model there is no explicit regulatory motive behind the existence of shadow banks or the market for securitised assets, although we proxy the advantage that shadow banks enjoy from being unregulated by allowing them to carry higher leverage than commercial banks. Relaxing some of these strong assumptions is a topic ripe for future work.

News and labour market dynamics in the data and in matching models

Summary of Working Paper No. 488 Konstantinos Theodoridis and Francesco Zanetti

A number of studies establish that anticipated changes in future productivity, referred to as news shocks, represent an important source of business cycle fluctuations. Many authors have focused on the effect of news on economic activity, but none so far have investigated the effect on labour market variables. This paper fills this gap. It develops a multivariate statistical model that identifies the effect of anticipated productivity shocks on unemployment, wages and the job finding probability, and it then investigates to what extent a simple theoretical model with real frictions on the labour market is able to replicate the empirical impact of news shocks on labour market variables and macroeconomic aggregates.

In the aftermath of a positive news shock, unemployment falls, whereas wages and the job finding probability increase. The inclusion of labour market variables does not alter the response of macroeconomic aggregates to the news shock, since output and investment modestly fall and consumption increases, in line with recent studies that abstract from labour market variables. We establish, using US data, that news shocks explain 30% of unemployment fluctuations and approximately 20% of the job finding rate, whereas their contribution to output and consumption is more limited to around 15% in the long run. We also illustrate that most of the historical fluctuations in the job finding rate and unemployment are explained by news shocks, whereas news shocks play a limited role in explaining wages and output fluctuations.

We next set up a simple theoretical model with real (search and matching) distortions in the labour market. We find that this basic framework replicates the news shocks identified in the data relatively well. The theoretical model shows that in response to a positive news shock the firm anticipates that the surplus from establishing a match increases, thereby leading to an increase in vacancy posting that generates a decrease in unemployment. High vacancy posting and low unemployment raise labour market tightness, which increases the job finding rate. In general, the qualitative responses are similar to those from the time-series model. However, the responses of unemployment and wages to a news shock are not in line with those predicted by the empirical model. Hence, we investigate to what extent refinements to the basic framework improve the model's performance. We establish that the job destruction rate and real wage rigidities are important for the response of unemployment and wages to the news shock and for the overall variables' responses.

Expectations, risk premia and information spanning in dynamic term structure model estimation

Summary of Working Paper No. 489 Rodrigo Guimarães

Market interest rates are of great interest to policymakers, not least because they play a crucial role in the monetary transmission mechanism. Moreover, financial market measures of future interest rates and inflation rates can also provide useful and timely information when making policy decisions.

This information complements and extends other sources monitored by policymakers, such as surveys of private forecasters and macroeconomic forecasting models. Market rates are available at a much higher frequency and for longer forward horizons than other data, as well as being available in a long time series. This can prove crucial in answering questions that involve the reaction to policy (such as announcements), comparisons over long periods (the effect of institutional changes, such as independence of the central bank), or effects that are expected to have distinct effects over different horizons (such as forward guidance).

In order to extract policy-relevant information from yields, it is important to understand what has driven these rates lower. Decompositions can be carried out along a number of dimensions to shed light on the drivers. First, movements in interest rates can be split into movements at different forward horizons to assess whether the changes are mainly at shorter or longer horizons. Second, movements in nominal rates can be decomposed into changes in real interest rates and changes in implied inflation rates.

And third, movements in market rates can be decomposed into two parts; one that reflects changes in expectations of future short-term rates, and another associated with changes in their required compensation for risk ('risk premia'). Disentangling both is important for policymakers because influencing the expected path of the policy rate plays an important role in the transmission mechanism of monetary policy. And estimating risk premia can give policymakers an indication of market participants' assessments of the perceived risks. In addition, some measures are designed to reduce the compensation for risk (such as quantitative easing).

While the first two decompositions — time horizon and the real versus inflation split — can be done using available data, the distinction between expectation and risk compensation components is more complicated. Extracting this information requires complicated theoretical models and statistical techniques, which raises the question of reliable decompositions. Unfortunately, the most popular class of models, both within academia and with major policy institutions, are known to be subject to instability problems that would hamper their use for policy. This paper focuses on how to obtain robust estimates from these models for the quantities of interest for policymakers: the expected path of future interest and inflation rates as well as real and inflation risk premia.

We analyse the robustness of the decomposition obtained from the workhorse model in previous work, the family of Gaussian affine dynamic term structure models. The great advantage of this type of model, which assumes linearity and a relatively straightforward probability distribution of shocks to returns, is its tractability. At the same time their flexibility is a great asset, necessary to accommodate the rich behaviour of bond yields observed over time and across maturities. But without enough restrictions or information to pin down the model parameters, this flexibility can become a liability, resulting in instability in the implied decomposition into risk and expectations.

Exploring recent advances in yield curve modelling this paper compares alternative methods proposed in previous work to ensure sensible decompositions. These include using survey forecasts from professional forecasters, restricting the way risk premia are allowed to vary or purely statistical techniques. This paper finds that using surveys of private professional forecasters to help anchor the model dynamics is the most reliable way to obtain robust decompositions.

In addition, the use of surveys automatically delivers 'sensible' decompositions because these survey forecasts (i) have been shown to provide good proxies for expected future rates (good forecasting properties); (ii) are true real-time measures (not subject to look-ahead or overfitting biases); (iii) can incorporate information that is readily available to practitioners (political events, changes in policy or policy frameworks) hard to obtain from past data.

The outputs of the models with surveys have been used to analyse the evolution of UK government bond yields in a 2012 Q3 *Quarterly Bulletin* article. The model decomposition of nominal, real and market inflation rates provided valuable insight about the behaviour of yields. It proved particularly useful in understanding the recent period of the financial crises and how it impacted market rates.

In a more technical contribution, the paper also links the ability of surveys to stabilise the decomposition of yields to new developments in term structure modelling related to spanning of information. A Monte Carlo study (based on random simulations of a theoretical model) confirms the importance of having additional information about future dynamics to reliably estimate these models. It suggests that the introduction of surveys delivers gains in precision equivalent to observing at least twice as long a sample — in other words we would need double the amount of information available (wait another 40 years) to obtain measures as reliable as those we can obtain by adding surveys.

Adaptive forecasting in the presence of recent and ongoing structural change

Summary of Working Paper No. 490 Liudas Giraitis, George Kapetanios and Simon Price

Forecasting is an important activity for central banks, not least because policy takes effect with a lag. Inevitably, policy is forward looking. Thus in many central banks, including the Bank of England, the published forecast is a key tool in communicating judgements about monetary policy and the economy. The Bank's forecast, published in the *Inflation Report*, represents the judgements of the Monetary Policy Committee and is not mechanically produced by a single model. However, many forecasting models — a 'suite' of models — help the Committee determine its judgement, including simple largely atheoretical models of the type considered in this paper.

One common cause of forecast failure is that structural changes or 'breaks' keep on occurring in the underlying relationships in the economy, and this paper addresses that problem, building on previous work undertaken in the Bank. The problem, almost by definition, is that we do not know what form the structural break took. If we did, we could model it: but then it would not be a structural break, but a known data-generation process. What we need are methods that are useful where there is the possibility of a wide range of types of structural change. The earlier work showed that a robust way of forecasting in such an environment is to discount past data so that more recent data is given more weight. This helps avoid forecast errors, as if there have been structural breaks in the past, the data pertaining to that period is given less weight compared to recent data where there may have been no or fewer breaks. This can be done in many ways. These include 'rolling windows' where all data before a cut-off date is excluded, exponentially declining weights smoothly lowering the weight for distant data (often implemented as an exponentially weighted moving average), and other methods. But this raises the practical question of exactly how rapidly to downweight. The innovation of the paper is to choose this by using in-sample forecast performance.

The paper shows that in a wide variety of situations the method will have good statistical properties. What is more, it can handle any degree of persistence. Speaking somewhat loosely, 'persistence' is the tendency for a series to be affected by its past behaviour. For example, a series that is simply a constant with some random white noise has no persistence. (In this case, the best forecast is to use all the data to calculate the mean as precisely as possible.) By contrast, in the classic random walk a series is equal to what it was last period plus a random white noise error, and so there is a high degree of persistence. (In that case, the best forecast ignores all except the last observation.) These examples show that the optimal rate of discounting past data is likely to depend upon persistence. We are also able to demonstrate that the method is very flexible. There are ways of including dynamics, similar to the widely used autoregressive (AR) method, known to produce good forecasts, where the series is solely related to a few of its own lags. We can also allow the weights to vary very flexibly using a non-parametric method which does not tie down the model to a specific form, and allow for other explanatory variables. The theory is for large samples, but we show using simulation ('Monte Carlo') methods that the methods work for short samples as well.

The proof of the forecast pudding is in the testing, so we apply the methods to a large number of economic variables from the United States using a sample from 1960 to 2008, comparing root mean square forecast errors, which is a standard criteria that penalises large forecast errors. Not all the series exhibit breaks, but in the typical (median) case the methods do better than an AR benchmark. The methods that work best are ones that allow for some dynamics. For some variables, such as financial spreads and some inflation series, they do spectacularly better. Moreover, in many cases the methods are significantly better (in the statistical sense) than the benchmark, meaning that they do better much more often than would be expected by chance.

We conclude that the proposed technique of downweighting past data in a way determined by past forecast performance is likely to be a useful item in the forecaster's toolkit.

Household debt and the dynamic effects of income tax changes

Summary of Working Paper No. 491 James Cloyne and Paolo Surico

The persistent rise in mortgage debt across many industrialised economies prior to the recent financial crisis has drawn considerable attention to the role of private indebtedness in the transmission of macroeconomic shocks.

A common presumption behind many existing studies on the relationship between debt and consumer behaviour is that debtors are more likely to face liquidity constraints and thus adjust their consumption significantly in response to conditions that unexpectedly change their income. An important implication is that it is not wealth *per se* that determines the consumption reaction to income changes; households who made a large durable purchase — such as housing — may well be wealthy and liquidity constrained at the same time, depending on their level of indebtedness as well as their mortgage repayments relative to earnings.

Despite the clear relevance of this transmission channel, little is known about whether household consumption behaviour varies with their debt position and whether the mortgagors' reaction to income changes is larger or smaller than outright homeowners'. At least three considerations make this task particularly challenging. First, survey data with good expenditure coverage typically lack equally detailed and reliable information on the household finance position over a sufficiently long period of time. Furthermore, whether a household holds mortgage debt or not is partly a choice. Last but not least, consumption and income changes are jointly determined so that one ideally needs to isolate a surprise movement in income.

To address the endogeneity of income changes, we exploit variation in aggregate income taxes. The United Kingdom is a natural choice for our purposes because there have been a large number of income tax changes in the past 40 years. Furthermore, detailed information from official documents allows us to identify individual tax measures and their motivation. Tax changes that were introduced for reasons unrelated to the business cycle can then be used to identify exogenous variation in household income.

To elicit individual debt positions, we group households by their housing tenure, which allows us to distinguish between homeowners with a mortgage and homeowners without. A further advantage of looking at heterogeneity in consumption through housing tenure is that we can investigate the dynamic effects of tax changes on another interesting group of British households, namely those renting from local authorities or housing associations: 'social renters'. A typical household in this group is characterised by little (if any) net wealth, low income and only compulsory education. These features fit well the traditional stereotype of liquidity-constrained households in one-asset models.

Using a long span of household survey data from the Family Expenditure Survey and a new narrative measure of tax shocks, our disaggregated approach allows us to identify the expenditure response of different groups of households to an income change. We find that the estimated dynamic effects are highly heterogeneous across housing tenures, in a way that is understated — if not missed — when households are grouped by age and/or education.

We establish a number of specific results. First, mortgagors exhibit the largest and most significant response. In contrast, outright homeowners hardly adjust their expenditure, with effects that are never statistically different from zero. Second, the response of social renters is significant but smaller than (though rarely statistically different from) the mortgagors'. Third, the composition of mortgagors' net wealth is significantly different from those of outright owners and social renters. More specifically, we show that a typical household with mortgage debt holds very little *liquid* net wealth despite owning sizable illiquid assets.

Liquidity constraints for lower income/less educated households is an often cited explanation for the aggregate effects of tax changes on GDP and consumption that are typically reported in the empirical macro literature. On one hand, we show that social renters do respond to these changes in their income. This type of (arguably) liquidity-constrained household, however, accounts for only around 20% of our sample and therefore appears to constitute too small a share to drive the aggregate results. On the other hand, mortgagors — who tend to hold little liquid wealth despite owning sizable housing equity — account for about half of the British population, thereby providing a novel interpretation for the aggregate effects of tax changes on the economy.

Generalised density forecast combinations

Summary of Working Paper No. 492 Nicholas Fawcett, George Kapetanios, James Mitchell and Simon Price

Forecasts have become central to monetary policy making. The Bank of England has for many years published forecasts of inflation and growth conditional on particular paths of the policy rate as a central element of the discussion of policy in the quarterly *Inflation Report*. For almost as long, the Bank has explicitly quantified the uncertainty surrounding the central forecast by publishing a forecast distribution or density around that path, in the famous 'fan charts'. The Bank uses these as part of its communication strategy that help to explain the MPC's thinking. This process is influenced by many inputs, including various forecasts, and knowing how well those input forecasts perform is potentially helpful. More recently, the Bank has published probabilistic descriptions of other events, for example the probability of exceeding particular inflation rates or unemployment falling below a threshold specified under the Forward Guidance framework introduced in August 2013. In the area of financial stability, the probabilities of bank failure or financial crises occurring are of obvious interest. Clearly, forecast densities are key to the Bank's activities.

Recent research has revealed that some form of forecast combination is a very powerful method of improving forecast density accuracy, and this is the focus of the paper. One promising method is to choose fixed (linear) combination weights to maximise the log score of the combination. This object gives a high weight to forecast densities that assign a high probability to the events that actually occur. The difference between two average log scores is a measure of 'distance' between two sets of density forecast, so we can compare densities. We generalise the simple linear approach by letting the combination weights follow general schemes. Specifically, we let the combination weights depend on the variable one is trying to forecast, by allowing the weights to depend on which specific interval the forecast lies in. This allows for the possibility that while one model may be particularly useful (and receive a high weight in the combination) when the economy or market is in recession or a bear market, for example, another model may be more informative when output growth is positive or there is a bull market. It also allows some time variation in the weights as forecast outcomes move into different regions of the forecast density.

We show that these generalised density combinations or pools have good theoretical properties. There are practical issues to be resolved. The number of intervals has to be decided: we determine this by a 'grid' search over a range of possibilities, choosing the one that maximises the log score in a part of the data sample held back for this purpose (a method known as cross-validation). Similarly, the specific values delimiting the intervals must also be estimated.

Theoretical properties are good, but in practice data are limited and not well behaved. So we try a range of 'Monte Carlo' simulations where we simulate different processes and sample sizes to see how the method might work in practice. These show that the generalised combinations are more flexible than their linear counterparts and in general can better mimic a range of true but unknown densities, irrespective of their forms. Although this additional flexibility does come at the expense of the introduction of additional (interval) parameters, the simulations indicate that the benefits of generalised combinations mostly survive the extra parameter estimation uncertainty; and increasingly so for larger sample sizes and more distinct component densities. As a practical matter, the number of intervals chosen is generally quite small. One experiment uses densities known to work well for inflation. Here too the generalised method outperforms the linear one, showing that the method is likely to be useful in macroeconomic applications. We also show how the generalised combinations can work better in practice, finding that they deliver much more accurate density forecasts of the S&P 500 daily returns than optimal linear combinations of the sort used in earlier work by researchers on this topic. Not only are the gains large in absolute terms, but they are extremely significant in the statistical sense.

In summary, forecast densities are important, and a good way to generate them is by combining individual densities. We generalise a method known to work well, and we find that in theory and in practice the gains are very large. Although we only apply the method to stock market returns where there are many thousands of observations, the Monte Carlo experiments suggest it will be effective in macroeconomic cases too.

The macroeconomic effects of monetary policy: a new measure for the United Kingdom

Summary of Working Paper No. 493 James Cloyne and Patrick Hürtgen

The precise impact of monetary policy on the macroeconomy is of obvious interest to central bank policymakers. Yet despite considerable research in the academic literature, there remains disagreement about the effects. A range of empirical estimates have emerged in the literature, and the effects on prices and output of a 1 percentage point innovation to the policy rate tend to be between 0.5% and 1%. But a notable exception is the so-called narrative method pioneered by Christina and David Romer in 2004, which found considerably larger effects. To our knowledge, and despite the attention given to these results, there are no other applications of this methodology to isolate interest rate changes that can be used to estimate the effects of policy. In addition, much of the empirical research on monetary policy has focused on the United States and there are far fewer results for other countries such as the United Kingdom. This paper fills both these gaps, providing new narrative-based estimates of the effect of monetary policy in the United Kingdom.

We focus on the effect of changes in the central bank's policy interest rate rather than on unconventional policy measures such as quantitative easing. While the effect of unconventional measures is clearly an important topic in its own right, interest rates still remain a key policy instrument. Furthermore, looking at changes in policy interest rates is important for understanding the effects of monetary policy in the past and to be comparable with the existing literature. The effect of interest rates on the macroeconomy therefore remains of considerable interest, both to macroeconomists and policymakers.

Identifying the effects of changes in monetary policy requires confronting at least three technical challenges. First, monetary policy instruments, interest rates, and other macroeconomic variables are determined simultaneously as policymakers both respond to macroeconomic fluctuations and intend their decisions to affect the economy. Second, policymakers are likely to react to expected future economic conditions as well as current and past information. Third, policymakers base their decisions on 'real-time' data (that available at the time), not the *ex-post* (revised) data often used in empirical studies.

A major advantage of the Romer and Romer approach is that we can directly tackle all three of these empirical challenges. First, we need to disentangle cyclical movements in short-term market interest rates from policymakers' intended changes in the policy target rate. A particular advantage of studying the United Kingdom is that the Bank of England's policy rate, Bank Rate, is the intended policy target rate. We therefore do not need to construct the implied policy target rate from central bank minutes as Romer and Romer did. As a second step, the target rate series is purged of discretionary policy changes that were responding to information about changes in the macroeconomy. This may include real-time data and forecasts that determine the policy reaction to anticipated economic conditions. We use historical sources to reconstruct a proxy for the information available to policymakers at the time. Specifically we construct an extensive new data set of historical Bank of England forecasts, private sector forecasts from the National Institute of Economic and Social Research and real-time data, for the sample from 1975 to 2007. (We hope that our detailed new data set and policy change series will provide a useful resource in itself.)

We perform a first-stage estimation to purge the intended policy target rate of systematic policy changes, producing a new series of policy changes that can be used to estimate the effects on the economy. Armed with our new measure of monetary policy changes we estimate the effects of monetary policy on the macroeconomy. In our baseline empirical specification, a 100 basis point tightening leads to a maximum decline in output of 0.6% and a fall in inflation of 1.0 percentage points after two to three years. Monetary policy changes have a protracted effect on the economy. Our results also suggest that GDP responds by a comparable magnitude to industrial production — around 0.5% at the peak.

The narrative results for the United States generated considerable discussion given the large effects found. Also employing this narrative approach, we find similar effects of monetary policy for the United Kingdom and the United States. Furthermore, we show that differences across empirical methods used in the literature, at least for the United Kingdom, largely result from differences in the implied paths for policy. Once we control for these effects we find our estimates are in line with the magnitudes reported elsewhere in the wider literature.

However, unlike many studies in the literature, and in keeping with Romer and Romer for the United States, we find a negative, significant and theoretically plausible response for inflation and prices. Our approach therefore solves the so-called 'price puzzle' — first documented by Christopher Sims in 1992 — for the United Kingdom, where prices and inflation puzzlingly increase following a monetary contraction. Investigating the issue further, we find that use of our new forecast data set is crucial for this result.

The effect of changes in monetary policy continues to be keenly debated, both in academic and policy circles. Furthermore, it seems likely that interest rates will remain a key tool in the future. Our estimates therefore contribute new evidence to the empirical literature. In doing so, we provide a rich new data set which we hope will provide exciting scope for future research.

Estimating the impact of changes in aggregate bank capital requirements during an upswing

Summary of Working Paper No. 494 Joseph Noss and Priscilla Toffano

This paper estimates the effect of changes in capital requirements applied to all UK-resident banks' aggregate capital requirements on lending during a credit boom. It is a 'top-down' study that investigates the joint dynamics of the aggregate capital ratio of UK-resident banks and a set of macrofinancial variables, including lending growth. Its results may be useful to policymakers given the growing international consensus on the need to apply time-varying macroprudential bank capital requirements on top of existing microprudential requirements. An example is the countercyclical capital buffer that comprises part of the toolkit of the Bank of England's Financial Policy Committee.

Estimating the effect of an increase in aggregate bank capital requirements on the macroeconomy is complicated by how such a policy tool has never before been used. There are, moreover, very few changes to aggregate regulatory capital requirements observable in past data. And for those changes in regulatory capital requirements that have occurred, it is difficult to isolate how much of the change in bank lending behaviour was as a result of those changes, rather than broader macroeconomic developments affecting banks. The approach offered here surmounts this problem by identifying shocks in past data that match a set of assumed directional responses of other variables to future changes in banks' aggregate capital requirements.

This analysis estimates how an increase in macroprudential capital requirements might affect banks' lending in the face a credit boom. In doing so, it assumes that an increase in banks' capital requirements have a negative effect on the supply of bank lending, at least in the short run. It is also important to note, however, that this assumption is likely to hold true only during a boom in the extension of credit, such as that witnessed before the recent financial crisis. It may not match the response of banks to regulation after the crisis, when, for example, an increase in macroprudential capital levels could improve investor confidence in the health of banks, allowing their cost of funding to fall, and thus enabling them to increase their level of capital without decreasing their lending.

The estimates of the impact of aggregate capital requirements on lending may — in certain states of the economic cycle — provide policymakers with a plausible 'upper bound' on the short-term effects of future increases in macroprudential capital requirements. This analysis concludes that an increase of 15 basis points (one standard deviation) in the aggregate capital ratio of the UK banking system is associated with a median reduction of around 1.4 percentage points in the level of lending after 16 quarters. The effect is found to be larger on total bank lending to corporates, and less on that to households, perhaps reflecting differences in capital requirements on lending to each sector. The impact on GDP growth is statistically insignificant.

The productivity puzzle: a firm-level investigation into employment behaviour and resource allocation over the crisis

Summary of Working Paper No. 495 Alina Barnett, Adrian Chiu, Jeremy Franklin and María Sebastiá-Barriel

Labour productivity growth in the United Kingdom has been exceptionally weak since the 2007/08 financial crisis and currently lies around 14 percentage points below the level implied by its pre-crisis trend growth rate. This phenomenon is commonly referred to as the UK 'productivity puzzle'. Such a prolonged period of weakness in labour productivity stands out from historical and international experiences.

This paper uses firm-level data collected by the Office for National Statistics (ONS) to improve our understanding about the drivers of the weakness in UK labour productivity. This analysis only covers the period to 2011, so cannot shed light on the strength in UK employment in 2012 and 2013.

There are two main objectives. First, to set out some stylised facts about productivity across both time and firm dimensions. Within this part we also discuss how representative our results are for the UK economy as a whole by comparing employment and gross value added in our sample of firms to ONS aggregate statistics. We observe that the productivity experience across firms of different sizes has been varied. Labour productivity fell more sharply for small and medium-sized enterprises (SMEs) at the onset of the crisis than for large firms, but by 2011 the weakness in productivity relative to its pre-crisis trend appears relatively evenly distributed across SMEs and large firms. Second, we examine the changing role of resource reallocation on UK productivity growth before and after the financial crisis.

There are two key findings that stand out. First, we find that a large part of the persistent weakness in productivity can be accounted for by the fact that the proportion of firms with shrinking output and flat employment doubled from 11% in 2005–07 to 22% in 2011. At any point in the economic cycle we see some firms who are growing and hiring more workers, while other firms are shrinking and reducing employment. At the onset of the financial crisis there was a significant decline in the proportion of firms that were growing and hiring, and a rise in the number of firms that were shrinking and firing. But by 2011, a large concentration of firms emerged that had shrinking output but no change in employment. This suggests that firms were able to respond flexibly to weak demand conditions by retaining staff at the expense of measured productivity, suggestive of an opening up of spare capacity within firms. This result implies a more temporary or cyclical explanation of the productivity slowdown in the United Kingdom, as these firms may have been well placed to increase production quickly without hiring when demand conditions started to strengthen. However, the strength of recent hiring behaviour since 2012 implies that this may now be less of a factor.

The second key finding relates to the role resource allocation played in the slowdown of labour productivity growth. Reallocation is the process through which factors of production move from lower to higher productivity firms, helping to take advantage of market expansion opportunities and generate aggregate increases in productivity. We find that labour reallocation, which includes movements in labour brought about by company formation and dissolution, and within-firm productivity improvements were equally important in driving productivity growth between 2002 and 2007. However since the crisis, the role of reallocation fell significantly while the contribution of within-firm productivity to aggregate productivity growth turned negative. Taken together, we find that a third of the shortfall in labour productivity since the crisis can be attributed to slower resource allocation across firms. The rest is due to productivity weakness within firms. This result is indicative of a more persistent interpretation of the UK productivity slowdown. There may be a number of factors that have slowed the reallocation process; for example a disrupted financial sector or heightened uncertainty about the economic environment. Since our data are at the firm and not plant level, these within-firm changes will not include any reallocation effects across individual plants that sit within particular firm entities and may, therefore, underestimate the scale of these reallocation effects.

We believe that the effect of reallocation on measured productivity may have been low in part because a greater proportion of firms facing difficulties have managed to survive the latest recession compared to previous ones. As discussed in the Bank of England's August 2013 *Inflation Report*, the number of company liquidations has remained unusually low this recession given developments in output. There may be several factors that may have helped firms survive, for example the low level of Bank Rate, weak real wage growth and any forms of loan forbearance (the extent of SME forbearance is discussed in the 2013 Q4 *Quarterly Bulletin*). To try to understand the likely magnitude of this effect, we develop a scenario to examine what might have happened to measured labour productivity should firm deaths have increased to a level more consistent with the 1990s recession. Because our data set only goes back to 1997, and therefore does not include previous recessions, we use ONS aggregate statistics to inform our counterfactual exercise. We find that the unusually low level of business failure is likely to have materially lowered measured labour productivity. Nonetheless, lower business failure, and the resultant lower unemployment, probably meant that the loss to GDP and general welfare associated with the financial crisis was smaller than it otherwise would have been.

Overall, our key findings suggest that the slowdown in UK productivity is likely to have coincided with an opening of spare capacity at the onset of the recession, as firms reacted flexibly to the weakness in demand by retaining staff. However, the strength in aggregate employment since 2012 implies that this may now be less of a factor. We also find that reallocation between firms (in terms of both the movement of labour and firm entry and exit) contributed significantly to aggregate productivity growth before the crisis, but its contribution fell substantially after. The speed at which labour productivity is able to grow in the short to medium term may be limited by the extent to which impaired allocation of resources across the economy continues to be a binding constraint. But exactly how companies and resources will respond as demand recovers remain key questions for the economic outlook.

Uncertainty in a model with credit frictions

Summary of Working Paper No. 496 Ambrogio Cesa-Bianchi and Emilio Fernandez-Corugedo

At around the time of Lehman Brothers' collapse, the world economy experienced a large contraction in economic activity. This was followed by an anaemic recovery and high levels of economic uncertainty by many measures. In the past few years, policymakers have often cited elevated uncertainty as a key reason why GDP growth in the United Kingdom and elsewhere has been sluggish. Despite this widely held belief that it has been important, the exact role that uncertainty plays in economic agents' decisions is hard to gauge. This makes it important to understand what the mechanisms may be, not least to be able to better offset any negative effects.

One long-held notion is that, when uncertainty about their future stream of income rises, households increase their savings and decrease their consumption. This is known as precautionary savings. But departures from the paradigm of frictionless financial markets can create additional channels through which uncertainty can affect macroeconomic outcomes. For example, when the relation between lender and borrower is subject to incomplete information, an increase in uncertainty will in general raise the cost of external finance.

This paper investigates the relationship between uncertainty and economic activity in a model that embodies both the precautionary savings and the cost of external finance channels. But what do we mean by uncertainty? There are many different notions and many different ways to model it. Economic theory generally hinges on the idea that macroeconomic fluctuations are the result of exogenous shocks to some key variables (for example to the level of productivity). In this paper we analyse the impact of uncertainty shocks, ie shocks to the variance rather than the level of those variables. For example, an uncertainty shock to productivity would increase its volatility without affecting its average level.

In our model, households consume, invest their savings in safe bank deposits, and supply labour to firms. In turn, firms produce goods with a technology whose productivity is subject to economy-wide shocks. Households' income depends on this aggregate productivity. Therefore, we define an increase in the volatility of aggregate productivity as a macro uncertainty shock.

Our model also has entrepreneurs who transform unfinished capital into finished capital with a technology that is subject to idiosyncratic (ie not aggregate but specific to the firm) productivity shocks. Entrepreneurs are heterogeneous, since their productivity will differ depending on the realisation of the idiosyncratic shock. And they need to borrow from banks to finance unfinished capital purchases. We assume that there are informational discrepancies between the lender and borrower which create a credit friction. We also assume that adverse shocks to idiosyncratic productivity may induce entrepreneurs to default on their debt. As a result, banks optimally charge a lending rate higher than the rate of return of a safe asset, thereby creating a spread between the two rates. An increase in the dispersion of entrepreneurs' productivity (or a micro uncertainty shock) induces banks to charge a higher spread, therefore reducing their demand for capital.

Our model simulations show that micro uncertainty shocks have a bigger impact on growth than macro uncertainty shocks. Specifically, we set the time-series properties of micro and macro uncertainty to match US data. We find that a two standard deviation increase in micro uncertainty generates a fall in GDP of 1.4 percentage points, relative to a fall of less than 0.1 in response to a two standard deviation shock to macro uncertainty.

Intuitively, while macro uncertainty operates primarily through the precautionary savings channel, micro uncertainty acts through the cost of external debt and capital demand and, therefore, it is greatly magnified by the credit friction.

Our results have important implications. Uncertainty shocks seem to have a first-order impact only when 'directly' interacted with credit market imperfections. We do not interpret this evidence as suggesting that uncertainty affects the economy mainly through investment and only to a lesser extent through consumption. Indeed, if households were to borrow in imperfect credit markets, the same amplification mechanism observed for entrepreneurs would be at work. Instead, our results suggest that uncertainty shocks can generate sizable impact on economic activity only when transmitted through a credit channel. It may be that when financial markets are performing better, the impact of uncertainty may diminish.

The international transmission of bank capital requirements: evidence from the United Kingdom

Summary of Working Paper No. 497 Shekhar Aiyar, Charles W Calomiris, John Hooley, Yevgeniya Korniyenko and Tomasz Wieladek

The global financial crisis of 2008 has led to an increasing focus on macroprudential regulation. An important element of macroprudential regimes going forward will be time-varying minimum capital requirements on banks. Higher capital requirements could make banks more resilient to adverse shocks. But as part of their adjustment to higher capital requirements, banks may cut back lending, especially if capital requirements are binding and bank equity is more costly to raise than bank debt. In light of these theoretical predictions, previous work has examined the impact of changes in capital requirements on the domestic loan supply.

But there is little reason to think that the response to capital requirement changes would be restricted to the country in which the regulatory change originates. In this paper we examine whether a rise in microprudential minimum capital requirements on UK banks is transmitted to foreign economies through a change in the supply of cross-border credit for the period 1999 Q1–2006 Q4. The United Kingdom provides an ideal testing ground for this analysis, for at least two reasons. First, UK-resident banks tend to be very globalised, not just through affiliated banks abroad, but also through cross-border lending and liabilities. Second, during the 1990s and 2000s the UK microprudential regulator, the Financial Services Authority, imposed bank-specific, time-varying minimum capital requirements on the banks under its purview. Merging these regulatory data with detailed data on each bank's cross-border lending creates a unique database that is well suited to identifying the cross-border credit supply impact of minimum capital requirements. In particular, we can observe quarterly cross-border lending by each bank to up to 145 countries. The detailed recipient country-level data allow us to control for demand with fixed effects and therefore give a loan supply interpretation to our estimates.

We find that a change in minimum capital requirements indeed elicits a robust cross-border supply response by affected banks: a 100 basis point increase in the capital requirement is associated with a reduction in the growth rate of cross-border credit of 5.5 percentage points. Overall, this is broadly similar to the effects of between 5.7% and 7.6% reported in studies that focus on the transmission to the domestic credit supply. Banks also tend to favour their most important country relationships, so that the cross-border credit supply response in 'core' countries — defined as countries that tend to be important destinations for cross-border lending from the perspective of the individual bank — is significantly less than in others. Furthermore, we find that banks tend to cut back cross-border credit to other banks (including foreign affiliates) rather than to firms and households. That observation is consistent with a greater willingness, or ability, to cut back on shorter maturity, wholesale lending. This implies that an important part of the cross-border transmission of capital requirements occurs through a liquidity shock to foreign banking systems. We do not find a significant impact on direct cross-border credit to non-banks (ie firms and households).

The two faces of cross-border banking flows: an investigation into the links between global risk, arms-length funding and internal capital markets

Summary of Working Paper No. 498 Dennis Reinhardt and Steven J Riddiough

Cross-border funding between banks is an economically important source of finance. It is comprised of two distinctive forms of funding. First, there is *arms-length* (interbank) funding, that takes place between unrelated banks. Second, there is related (intragroup) funding that takes place between global banks and their foreign affiliates within an *internal capital market*. It has been documented how there is a risk that both forms of funding are withdrawn during periods of heightened risk in the global financial system, and economic theory predicts that the two forms should behave in the same way during a financial crisis.

Yet, the two forms of funding have key differences, which may mean they behave differently during a crisis. Within an internal capital market, a global parent bank has the power to shift liquidity from one part of its group to another. Additionally, a bank lending internally has more information about their counterparties' overall riskiness, relative to banks lending at arms-length. The differences could influence the way the two flows behave in response to fluctuations in risk in the global financial system. It is therefore possible that some countries' banking systems could be more insulated from heightened global risk than others, depending on their mix of interbank and intragroup funding and the share of intragroup funding held by global parent banks relative to foreign affiliates.

In this paper, we empirically study the behaviour of *disaggregated* cross-border bank-to-bank funding — disaggregated into interbank and intragroup funding — in relation to swings in risk in the global financial system. We do so by sequentially decomposing aggregate cross-border funding between banks, across 25 advanced and emerging market economies, using data on cross-border banking flows from the Bank for International Settlements. First, we split funding to banks in a particular country into two baskets (i) funding between *arms-length* counterparties (interbank flows) and (ii) funding between banks within the same banking group (intragroup flows). Next, to paint a more detailed picture, we further disaggregate intragroup funding between flows to parent banks and flows to foreign affiliate banks.

We find that a period of high and rising global risk aversion, such as that witnessed following the collapse of Lehman Brothers, results in markedly different behaviour in interbank and intragroup flows. Intragroup funding, which makes up around half of all cross-border funding between banks, *rises* when global risk increases and is invariant to periods of high global risk. Interbank funding displays the opposite behaviour — it is withdrawn during periods of high global risk, with emerging economies particularly vulnerable. These findings contradict the theoretical prediction that both interbank and intragroup flows will contract during periods of heightened global risk. In fact, each country's mix of interbank and intragroup funding alone, can explain up to 45% of the change in cross-border bank-to-bank funding across countries, following the collapse of Lehman Brothers. We also reveal further information about the behaviour of cross-border banking flows. For example, the decision to withdraw interbank funding during the financial crisis is found to have been closely related to whether a country was experiencing a systemic banking crisis.

We show that higher intragroup funding during periods of heightened risk is principally driven by global banks headquartered in advanced economies, receiving funding from their foreign affiliates. We find that banking systems with a high share of global banks were relatively well insulated against funding withdrawals during the global financial crisis. But we do not find evidence of significantly reduced intragroup funding to foreign affiliates in either advanced or emerging economies during periods of high global risk. In fact, we find that foreign affiliates resident in emerging economies experience an *increase* in intragroup funding, when the average profitability of banks in the local economy is low. This result is found to hold even during the financial crisis, indicative of the beneficial role financial globalisation can play for emerging economies with resident foreign banks.

Overall, the results call for policymakers and academics to focus attention on the disaggregation of cross-border bank-to-bank flows, as the contrasting behaviour of interbank and intragroup funding in response to fluctuations in global risk has implications for a banking system's financial stability.

Sectoral shocks and monetary policy in the United Kingdom

Summary of Working Paper No. 499 Huw Dixon, Jeremy Franklin and Stephen Millard

A key question for monetary policy makers is how to deal with 'relative price' shocks; that is, movements in individual prices that do not reflect aggregate inflationary pressure but that can, as a result of nominal rigidities, lead to temporary changes in inflation. This question has gained in importance in recent years as the United Kingdom has been affected by shocks to the price of food and energy, which fall into the category of relative price shocks. In order to get at this question, this paper develops a framework within which we can examine sectoral shocks, their effects on the UK economy and how monetary policy makers should respond to them. More specifically, our framework links together news in the consumer prices index (CPI) data at the sectoral level and the behaviour of the economy at the aggregate level. Such a framework can be used to address several questions about the links between prices at the aggregate and sectoral levels as well as the particular issue of how monetary policy should respond to movements in sectoral prices.

Before constructing our model, we first investigate the empirical properties of quarterly sectoral inflation rates in the United Kingdom over the period 1988–2011. The idea is to generate some stylised facts with which we would like our model to be consistent. We find that the sectoral rates have much bigger variances than aggregate CPI inflation and that there is little cross-correlation of inflation across sectors. We also find that the persistence we observe in aggregate inflation comes mainly from the effect of the aggregate factors with sectoral shocks being white noise. Leaving aside food and energy, we find that sectoral shocks explain the majority of the variance of sectoral inflation rates, but explain little of the variance of aggregate inflation, which is mostly explained by macroeconomic factors.

We then use the UK CPI microdata for the period 1996–2006 to calibrate a 'Generalized Taylor Economy' (GTE) for each of the twelve Classification Of Individual Consumption by Purpose (COICOP) sectors. The idea of a GTE is that price changes are staggered with some firms in the sector changing their prices every quarter, some changing their prices every two quarters, and so on up to some who only change their price every twelve quarters. To calibrate the proportion of firms who change their prices every so many quarters, we estimate the cross-sectional distribution of firms whose prices have different durations within each COICOP sector. We can then use our model to trace out the effects of a productivity increase or decrease affecting a particular COICOP sector. The GTE model of pricing is then embedded into an open-economy macroeconomic model of the United Kingdom in which we separate food and energy out of the CPI sectors, giving them an independent role. We do this because these are both sectors where prices are largely determined outside the United Kingdom and have had a significant impact on inflation in specific periods. Doing so, in turn, enables us to examine the issue of the extent to which monetary policy should or should not respond to movements in food and energy prices.

The policy issue on which we focus is how monetary policy should respond to sectoral shocks, or whether it should concentrate instead on some measure of underlying inflation. In this paper, we look at two such measures: one that strips out the most volatile components of CPI inflation from the index and a second that strips out that part of CPI inflation that can be thought of as being 'external' to the United Kingdom, leaving only 'domestically generated' inflation.

In our model, we look at simple rules in which the central bank alters interest rates in response to movements in aggregate and sectoral inflation rates and output relative to trend. We find that the optimal rule in which interest rates respond to sectoral inflation rates leads to a small improvement over a rule in which interest rates only respond to aggregate inflation. However, this gain comes from partially looking through movements in aggregate inflation driven by movements in petrol price inflation, which is volatile and tends not to reflect underlying inflationary pressure.

Modelling the service sector

Summary of Working Paper No. 500 Philip King and Stephen Millard

In this paper, we try to understand better how service sector companies operate and to incorporate some of these features into an otherwise standard macroeconomic model so as to examine their implications. We have two motivations for doing this. First, in the wake of the financial crisis output fell dramatically while inflation remained above its target and productivity collapsed relative to its previous trend. The fall in productivity relative to trend was pronounced within the service sector, and then most particularly in certain subsectors such as 'Professional, Scientific and Technical Activities'. At the same time, CPI services inflation has remained in the 3% to 5% corridor it has occupied since at least 2000. Given the weight of services in the economy — 75% in GDP and 50% in the CPI — it would seem that understanding how this sector works is crucial if we are to understand how the economy as a whole responds to shocks. Second, most standard macroeconomic models assume that 'value added' is produced using capital and labour and raw materials and imports are combined with 'value added' to produce final output. Unfortunately, this model is not particularly representative of what happens in the service sector. For example, how do we measure the real output of, say, a firm of consultants, architects or estate agents? And what are the inputs of such firms? It is clear, for instance, that human capital and other forms of intangible capital such as goodwill, firm-specific knowledge and ways of doing things, and client bases, to name but a few, will be extremely important in enabling service companies to produce output. And these factors are also likely to affect price and wage-setting in the service sector. For example, given the difficulty in measuring output and hence productivity, together with the importance of individual-specific human capital, how do you determine wages in a service company?

In order to get a better idea of how service sector firms actually operate in practice, we first embarked on a series of structured visits to a set of firms that span the service sector. More specifically, we visited around 30 private sector service providers, with a roughly even spread across Standard Industrial Classification sectors. In each case, we asked the firm what they considered to be their outputs and inputs and how they went about measuring them; we asked them what they considered to be full capacity and how they might respond to increases in demand; and we asked them about the form that their investment undertook and, more generally, about how they were able to achieve improvements in productivity. Our visits suggested two important features of service sector firms: the need to spend time on 'marketing' given the search and matching frictions present in the market for, in particular, business services, and the high degree of 'scalability' of many services.

Armed with these insights about how the services sector works, we wanted to understand the macroeconomic implications. We therefore incorporated these features into an otherwise standard macroeconomic model and examined the response of output, inflation and sectoral and aggregate productivity to sector-specific productivity shocks and aggregate demand shocks. Our results suggest that, in sectors where these features were important, productivity should respond negatively to negative demand shocks. This contrasts with the positive response of productivity to negative demand shocks in standard models. In that case, the capital stock takes time to adjust, so when demand declines, employment declines more than output and productivity goes up.

We then used the model to examine the effect of a negative demand shock caused by a rise in spreads in line with that seen during the financial crisis. We find that the model can explain a small but important part of the observed fall in business services productivity, and a small but less significant part of the fall in productivity in 'scalable' services. We feel that our approach to modelling services has been successful given that we have matched qualitatively the fact that business services productivity has performed particularly badly since 2007 and the anecdotal evidence that this has been associated with an increased proportion of the workforce in these companies used in tasks such as winning and maintaining contracts and trying to build up customer relationships more broadly. We conclude that it is important to incorporate these features into our macroeconomic models if we are to understand the evolution of economies in which the service sector is so important, such as the United Kingdom.