Likelihood inference in non-linear term structure models: the importance of the lower bound

Summary of Working Paper No. 481 Martin Andreasen and Andrew Meldrum

The yields on government bonds are of interest to monetary policy makers partly because they reflect financial market participants' expectations of future policy rates. As with any asset price, however, they also reflect the additional return — or 'risk premia' — that investors require to compensate them for the uncertainty surrounding future returns on the asset. And yields also play an independent and important role in the transmission mechanism of monetary policy. Central banks therefore make widespread use of models to both forecast yields and to decompose them into expectations of future policy rates and risk premia.

Perhaps the most popular type of model among central bankers, academics and financial market practitioners is the 'affine term structure model' (ATSM), where yields are a linear function of some underlying variables. This makes for tractability. These statistical models of bond yields are consistent with the standard assumption that investors cannot make risk-free arbitrage profits (ie investors cannot make profits by buying and selling different categories of bonds in such a way that the expected return from holding that portfolio is positive). But ATSMs do not impose the restriction that nominal interest rates are subject to a lower bound. This feature of the model is likely to have become more important in recent years given the historically low level of nominal bond yields.

Quadratic term structure models (QTSMs), in contrast, are more general and can be specified to be consistent with a lower bound. They are, however, substantially harder to estimate than ATSMs. This paper demonstrates for the first time that it is possible to use a numerical technique known as 'Particle Markov chain Monte Carlo' to estimate these models. This technique involves the random generation of many different candidate values for the model parameters. Each candidate draw of parameter values depends on the previous draws. Whether the candidate is accepted or rejected depends in part on how well it matches the observed data. This in turn is established using a different simulation technique known as a 'particle filter', which involves simulating many possible scenarios from the model and establishing how likely each scenario is given the observed data. Once we have considered a sufficiently large number of draws, the distribution of possible parameters will cease to change, known as convergence. This way of estimating these models has some desirable features relative to the methods that have been used previously. In particular, the statistical properties of the estimated model parameters can be more accurately established.

We apply the technique to estimate a QTSM using US nominal bond yields for the period 1962–2012. We find that the presence of the zero lower bound on nominal interest rates has important implications when using term structure models to forecast bond yields and short-term policy interest rates. Standard ATSMs imply around a 5%–15% probability of negative policy rates in ten years' time throughout the estimation period. During the recent financial crisis the ATSM implies probabilities of negative policy rates of more than 40% at shorter horizons. The QTSM rules this out by construction. The difference between policy rate forecasts from the two models becomes more important as bond yields approach the lower bound.

Has weak lending and activity in the United Kingdom been driven by credit supply shocks?

Summary of Working Paper No. 482 Alina Barnett and Ryland Thomas

The recent financial crisis has focused attention on the importance of credit supply and other financial shocks on the real economy. Prior to the crisis macroeconomists were typically interested in explaining movements in macroeconomic variables in terms of only a small number of aggregate level shocks, such as those from aggregate supply, aggregate demand and monetary policy. As a result the specific role of credit and financial market shocks were implicitly subsumed within one or other of these aggregate macroeconomic shocks. This paper attempts to disentangle the impact of credit market shocks on lending and activity in the UK economy. In particular we address three related questions that have been prompted by the recent financial crisis:

(i) **Are shocks to the supply of credit more like aggregate demand or supply shocks?** There is a growing literature that suggests shocks to the credit market can have permanent effects on potential supply. In some models that can mean that inflation rises rather than falls in response to a contraction in credit supply and a fall in output. What does the UK evidence suggest?

(ii) **How does a credit market shock differ from a monetary policy shock?** Both have an observationally equivalent effect on loan rates in the economy, but are they similar enough that monetary policy is able to offset a substantial part of a shock to credit supply. And how easily can we distinguish their separate effects in the data? In particular, do credit supply shocks have an additional quantitative effect via rationing and other non-price terms in addition to an effect operating via loan rates?

(iii) **What has been the role of credit supply shocks in the recent crisis?** Have credit supply shocks rather than shocks that affect credit demand been the most important factor driving the slowdown of UK bank lending during the financial crisis? And how much of the slowdown in UK activity can we attribute specifically to UK-specific credit shocks and how much to other factors such as global activity and uncertainty?

To address these issues we estimate a structural vector autoregression (SVAR) model for the UK economy over a data set that goes back to the late 1960s. The SVAR approach involves estimating a set of variables where each variable is regressed on past movements of itself and the other variables in the system. The unexplained component of each variable is then decomposed into the impact of different fundamental or 'structural' shocks using a theoretically based set of sign and timing restrictions for the shocks we wish to identify.

In this paper we identify six structural shocks using this SVAR analysis. We use standard sign restrictions on the pattern of reactions on specific variables to identify the three standard macroeconomic shocks mentioned above that are typically analysed in this framework — aggregate demand, aggregate supply and monetary policy. These shocks are commonly identified as aggregate demand if it moves inflation and GDP in the same direction, whereas an aggregate supply shock moves them in the opposite directions. The sign of the interest rate impact is then used to distinguish between monetary policy shocks and other aggregate demand shocks. Namely, a monetary policy shock leads to output and interest rates moving in an opposite direction whereas other aggregate demand shocks would usually lead to an interest movement in the same direction as output (as monetary policy attempts to offset the impact). We then use an additional set of timing and sign restrictions to identify specific credit and financial market shocks. One of these is identified uniquely as a shock to the supply of credit by banks. The other two are identified as shocks to the corporate bond and equity markets that affect the demand for bank credit for a given level of activity in the economy. So overall we are able to identify a shock to credit supply and a number of shocks that will affect credit demand in the economy. The identified shocks look plausible when we use them to explain the past 50 years of UK economic history.

When we apply this analysis to the crisis we find that:

- **Credit supply shocks look more like aggregate supply than aggregate demand shocks.** Credit supply shocks that lower bank borrowing and output appear, if anything, to have a positive effect on inflation. Our analysis suggests that some of this may reflect an effect of credit supply shocks on the exchange rate as well as an effect on potential supply. This could reflect the importance of financial services in UK trade. That means that credit supply shocks are also significantly different to monetary policy shocks which push output and inflation in the same direction for a given impact on interest rates in the economy.

- **Credit supply shocks look to have an important quantitative dimension.** When compared to a monetary policy shock that has an equivalent effect on loan rates, the quantity of credit appears to move almost (three) times as much.

- **Credit supply shocks can account for most of the rise in credit spreads and most of the slowdown in bank lending over the crisis.** Shocks affecting credit demand only appear to be marginally important in 2010 and 2011.

- **Credit supply shocks can account for up to a half of the fall in UK GDP relative to its pre-crisis trend.** Other shocks to aggregate demand and supply appear to have also played an important role in driving weak demand. Monetary policy (both through interest rates and quantitative easing) appear to have had a significant role in offsetting these shocks.

Risk news shocks and the business cycle

Summary of Working Paper No. 483 Gabor Pinter, Konstantinos Theodoridis and Tony Yates

How does uncertainty affect the financial system and the aggregate behaviour of the economy? Recent events have led to increasing attention to the question of how uncertainty might shape the depth and duration of financial and economic crises. In addition, macroeconomists have emphasised the role of shocks originated in the financial system in driving macroeconomic fluctuations. This paper develops a multivariate statistical model as well as a theoretical framework to show that uncertainty related to financial markets has played a considerable role in explaining the past 30 years of US business cycles.

In our model, a financial disturbance is defined as an exogenous process that drives the dispersion of returns on investment. As these forces govern the state of investment risk in the economy, we refer to these perturbations as 'risk shocks'. Moreover, we distinguish between contemporaneous (unanticipated) and news-type (anticipated) components of these exogenous processes. By doing so, we build on recent academic papers which suggest that most of the economic effects of financial shocks occur as economic agents respond to advance information, 'news', about the future realisation of these processes. Some of these papers find that the overall effects of these disturbances to financial markets account for about 60% of output fluctuations in the United States.

The empirical part of our paper develops a multivariate statistical model which we use to identify risk and risk news shocks in the data. This allows us to quantify and distinguish the partial impact of risk and risk news shock from that of other, more standard, macroeconomic shocks such as monetary policy, supply and demand shocks.

Our empirical results suggest that the combined effects of risk and risk news shocks explain approximately 20% of US output fluctuations over the 1980–2010 period. This is a more modest effect than that found in previous studies. Nevertheless, we find that these types of financial disturbances have a large impact on the federal funds rate, suggesting that revelations about future uncertainty induce a vigorous and protracted response of the US monetary policy authority. With central bank rates pinned at their zero lower bound for some time now in the United States, United Kingdom and Japan, our results would suggest that risk news shocks may have impacted on the real economy more recently, and could in the future, until such time as conditions allow the central bank to raise rates to more normal levels.

The theoretical part of this study then develops a relatively standard quantitative 'dynamic stochastic general equilibrium' (DSGE) model. Models of this type capture the evolving and interconnected dynamics of the entire economy, allowing for the presence of random ('stochastic') shocks. The model is made realistic by the presence of various nominal and real frictions. These include the assumption that a fraction of households are 'non-Ricardian', meaning that they do not base their decisions on their expectations about future income, as they do not have access to financial markets and their consumption is a function of their current (rather than future) disposable income. In addition, our model features a form of 'financial accelerator' mechanism stemming from the riskiness of business loans in the model, as the returns on projects are subject to idiosyncratic (ie firm-specific) shocks. We refer to the distribution of these idiosyncratic shocks as risk shocks, reflecting on the underlying investment risk in our model economy. A sufficiently adverse draw from this distribution can make a particular borrowing firm insolvent, which causes lenders to charge an *ex ante* higher interest rate compared to the risk-free rate. This premium moves countercyclically with business equity (borrower's net worth) and procyclically with investment risk.

The estimated version of our theoretical model reveals that in order to match the quantitative responses of risk shocks implied by our statistical analysis, the degree of real rigidities in the model such as the fraction of non-Ricardian households must be remarkably high. From this, we conclude that there is still more work to be done in order to improve the endogenous propagation of financial shocks in DSGE models.

GDP-linked bonds and sovereign default

Summary of Working Paper No. 484 David Barr, Oliver Bush and Alex Pienkowski

This paper investigates the conditions under which GDP-linked bonds help to protect governments (or 'sovereigns') from unexpected poor growth outturns that might otherwise push them into a debt crisis. This is because the return on these bonds varies in proportion to the country's GDP — when growth is weak, the debt-servicing cost and repayment amount automatically declines; and when growth is strong, the return on the bond increases. This helps to stabilise a sovereign's debt to GDP ratio and makes it less likely that a deep recession will trigger a debt crisis and cause a default. GDP-linked bonds, therefore, can be viewed as a form of 'recession insurance' for sovereigns. While all countries might experience some benefit from the use of GDP-linked debt, economies with higher GDP growth volatility (such as emerging market economies) or countries where monetary policy is constrained (such as those in a monetary union) are likely to benefit most.

We use a calibrated model of sovereign default based on work by International Monetary Fund authors, which delivers a calculation of the maximum level of debt that a sovereign is likely to be able to sustain before it risks facing a crisis. This model is estimated for a 'representative' sovereign in two scenarios — (i) when all debt is issued as conventional bonds; and (ii) when all debt is in the form of GDP-linked bonds.

Given the simplicity of the model, these debt thresholds should not be interpreted as hard limits. In fact, historical experience suggests that many countries can exceed these levels without facing repayment problems. Instead, the focus of this paper is to consider how GDP-linked bonds can help to reduce the risks to a sovereign within this simple framework. This implies more attention should be focused on the amount GDP-linked bonds can potentially increase debt limits, rather than the absolute value of the debt limits themselves. Under the simplest model set-up we find that GDP-linked bonds have a substantial impact on a sovereign's debt limit — raising it by around 100% of GDP.

This analysis abstracts from a number of important considerations, so the baseline model is then made more realistic with two innovations. First, investors are now assumed to be risk-averse and require an additional premium to hold risky assets. This means that when the return on the asset is uncertain — either due to a risk of payment default, or in the case of GDP-linked bonds, because future growth outturns are uncertain — investors will charge a higher interest rate on debt. Second, it is assumed that when a sovereign changes its fiscal policy stance in order to try and stabilise debt, this has an impact on growth. For instance, when a sovereign increases its primary balance this will drag down on GDP growth. When these two modifications are included in the model, the additional 'fiscal space' derived from the introduction of GDP-linked bonds is around 45% of GDP.

The final section of analysis considers the welfare implications of issuing GDP-linked bonds. Sovereign defaults have the potential to damage the domestic economy significantly, so reducing the incidence of this will improve welfare. A stable and predictable fiscal policy is also desirable, as taxpayers are not faced with unexpected and erratic changes in tax. GDP-linked bonds help both to reduce the incidence of sovereign default and to stabilise fiscal policy. But on average taxpayers will have to pay higher interest payments on GDP-linked bonds (at least at low and moderate debt levels) compared to conventional bonds, which will lower taxpayer welfare. On balance, however, we conclude that GDP-linked bonds may provide a substantial net benefit in welfare terms — in our calibration this is equivalent to consumption equal to between 1% and 9% of GDP in perpetuity.

In summary, GDP-linked bonds have the potential to reduce the incidence of costly sovereign default and allow fiscal policy to be more stable and predictable. The welfare gains from this outweigh any additional costs associated with issuing such debt, especially for sovereigns with volatile GDP. GDP-linked bonds also have the potential to improve the functioning of the international monetary and financial system, by encouraging greater country self-insurance, and reducing the reliance on large-scale official sector support programmes to resolve crises.

Identifying channels of credit substitution when bank capital requirements are varied

Summary of Working Paper No. 485 Shekhar Aiyar, Charles W Calomiris and Tomasz Wieladek

One goal of macroprudential policy is to limit systemic risk by raising capital requirements in response to lending-fueled booms, whether at an economy-wide or sectoral level, so that banks will be able to weather adverse shocks from a sudden change in market conditions. The raising of capital requirements has two effects on financial resilience. First, it improves the capital position of banks. Second, to the extent that the capital requirement increase reduces the aggregate supply of credit, it may prevent credit-driven asset bubbles from forming in the first place. Given that a central channel of macroprudential regulation is the use of capital ratio requirements to control the aggregate supply of credit as a means of limiting systemic risk and maintaining financial resilience, policymakers need to gauge the extent to which changes in requirements on regulated banks affect the aggregate supply of credit. The two challenges in this task are identifying the effects of capital requirement changes on regulated banks and measuring the size of 'leakages' — the extent to which non-regulated forms of credit offset changes in the supply of credit from regulated institutions. This study explores the latter.

The size and nature of potential leakages, however, remains uncertain. In particular, leakages can occur through at least three different channels. First, a foreign-based banking group may operate both a foreign subsidiary, which is subject to UK capital regulation, and an affiliated foreign branch, which is subject to capital regulation in its home country. In that case, raising the capital requirement on the subsidiary may simply produce a shift of assets from the subsidiary to the branch. Second, interbank competition between domestically regulated banking enterprises and foreign branches operating in the same sectors of the economy can lead to credit substitution between the former and the latter. Finally, it is also possible that leakage occurs outside the banking system. Firms that experience reductions in bank credit may seek funding from capital markets.

The United Kingdom during the period 1998–2007 provides a unique environment for addressing highly policy-relevant questions about the nature of leakages as a result of changes in bank minimum capital requirements. The UK regulators set bank-specific capital requirements on the basis of perceived operational and market risks. Cross-sectional differences in capital requirements were large, and changes in bank-specific capital requirements were frequent. This paper focuses on identifying and comparing the relative strength of different channels of credit substitution in response to changes in banks' minimum capital requirements. We expect the substitutability of credit supply between regulated subsidiaries and affiliated branches to be greater than between regulated subsidiaries and unaffiliated branches, for several reasons. First, the affiliated branch has a stronger incentive to lend than an unaffiliated branch because it may be able to preserve a valuable lending relationship with relatively little effort on the part of loan officers. Second, the affiliated branch may be able to originate the loan at low transacting cost, by simply transferring the asset from one balance sheet to another. Finally, affiliated branch lenders would enjoy an information advantage about the impending change in regulatory policy toward the affiliated subsidiary. Changes in subsidiary capital requirements were not a matter of public information over our sample period. The affiliated branch would be privy to knowledge of the regulatory policy change affecting its affiliated subsidiary, and that information likely would be shared with the affiliated branch several weeks or months in advance of the change in the requirement.

Our results are consistent with these predictions. 'Leakage' between affiliated branches and subsidiaries is roughly twice as large as 'leakage' that arises as a result of interbank competition between UK-regulated entities and unaffiliated foreign branches competing in the same sectors of the economy. A simple calculation suggests both types of leakages together could offset aggregate changes in credit, following changes in capital requirements, by roughly 43.1%. But we do not find evidence for a reaction of securities issuance in response to changes in capital requirements. Regulators are of course aware of this problem. In particular, they have pledged to find ways to co-operate internationally to co-ordinate capital requirement policies in the interest of minimising leakage. Basel III contemplates a reciprocity arrangement whereby foreign regulators of branches located abroad will match changes in the host country's capital requirement over the cycle for this purpose.

In summary, our findings have significant implications for economic policy. They suggest that co-ordination among national regulators is important in ensuring that changes in capital requirements have the desired impact on a country's banking system. Current regulatory initiatives, such as Basel III and the European CRD IV directive, already attempt to address the problem of 'leakage' from foreign branches through a provision for international reciprocity. When the capital requirement in one country is raised, capital requirements on foreign branches operating in that country will be raised correspondingly by their home country regulator. By identifying and quantifying leakages from foreign branches, this paper validates the importance of the reciprocity component of the new regulatory framework.

The impact of capital requirements on bank lending

Summary of Working Paper No. 486 Jonathan Bridges, David Gregory, Mette Nielsen, Silvia Pezzini, Amar Radia and Marco Spaltro

This paper investigates the effect of changes in regulatory capital requirements on bank capital and lending to UK households and firms. It is an empirical study drawing on a new bank-by-bank data set, exploiting variation in individual bank capital requirements in the United Kingdom between 1990 and 2011. There are two key results. First, regulatory requirements impact bank capital ratios; banks typically rebuild the 'buffer' in their capital ratios above the regulatory minimum following an increase in that minimum requirement.(1) Second, changes in regulatory capital requirements affect bank lending. Results vary across sectors, but in response to an increase in capital requirements, loan growth typically falls in the year following the regulatory change and recovers within three years.

Empirical evidence on the link between regulatory capital requirements and bank lending is also of interest to policymakers. The financial crisis has led to support for the use of capital requirements as a tool to mitigate risks in the financial system. In the United Kingdom, the Financial Policy Committee (FPC) is responsible for setting time-varying capital requirements on sectoral lending.

The effect of such capital requirements might differ from the effect of microprudential policy. As a result, the results from our study cannot be directly mapped across to how changing capital requirements are likely to affect bank capital and lending in a macroprudential framework; but they provide a useful guide to how banks have adjusted their capital ratios and lending structure on average in response to past microprudential supervisory actions. For example, banks might take a different approach to restoring capital buffers when other banks are subject to the same policy change and measures are public; expectations of forthcoming policy changes might lead to earlier reactions by banks; and there might be a different degree of 'leakages' where entities not domestically regulated step in with new lending. Also, during the transition to higher global regulatory standards, increasing capital requirements might augment rather than reduce lending for initially undercapitalised banks if confidence effects boost their resilience and capacity to lend. Furthermore, macroprudential regulators are often required to consider the wider implications of changing capital requirements, which could include any adverse impact on lending — for example, while the FPC's primary objective is to protect and enhance the resilience of the UK financial system, it also has a secondary objective to support the economic policy of the Government.

This paper uses a rich new data set constructed at the bank group level. It matches high-quality lending data with supervisory data on bank capital and capital requirements. Supervisory data include confidential bank-specific and time-varying capital requirements set by the Bank of England and the Financial Services Authority (FSA) in the United Kingdom between 1990 and 2011, which allow us to estimate directly the relationship between changes in capital requirements and individual bank lending behaviour. Lending data are adjusted to give a unique measure of true lending flows, rather than relying on changes in stock positions as a proxy; and we analyse lending responses at the sectoral level, such that both credit supply and demand conditions are allowed to vary across different sectors of the economy.

The bank-by-bank data set is exploited using two sets of panel regressions. First, we regress the actual capital ratio held by each bank on that bank's regulatory minimum capital ratio. That allows an assessment of whether regulatory requirements affect the capital banks hold. Second, the loan growth of each bank to different parts of the economy is regressed on that bank's individual regulatory requirement and on its actual capital ratio. By estimating these two equations, both the direct impact of a change in capital requirements on lending and any indirect impact via the response of bank capital can be taken into account when plotting the response of bank lending over time.

These regressions suggest that changes in regulatory capital requirements did impact bank behaviour over the sample period. First, we find that changes in regulatory requirements typically lead to a change in actual capital ratios — in response to an increase in the minimum ratio, banks tend to gradually rebuild the buffers that they initially held above the regulatory minimum. Second, capital requirements affect lending with different responses in different sectors of the economy — in the year following an increase, banks tend to cut (in descending order) lending to commercial real estate, to other corporates and household secured lending. The response of unsecured household lending is close to zero over the first year as a whole. Loan growth mostly recovers within three years. Finally, preliminary analysis suggests that banks' responses vary depending on bank size, capital buffers held, the business cycle, and the direction of the change in capital requirements.

These findings contribute to the debate on whether the Modigliani-Miller propositions hold (ie whether changes in the composition of a bank's liabilities affect the bank's overall cost of funds and credit supply), in which case changing banks' capital requirements would not affect lending. In practice, the empirical literature has identified a range of frictions (with taxation of debt versus equity being frequently mentioned) such that the debt/capital structure of banks may not be neutral for credit supply. Our paper confirms that regulatory requirements tend to affect capital ratios permanently and credit supply temporarily.

(1) A bank's capital ratio is given by total regulatory capital as a proportion of total risk-weighted assets. A bank's capital 'buffer' is given by the actual capital ratio minus that bank's minimum required capital ratio, as determined by the regulator.