

UK deposit-taker responses to the financial crisis: what are the lessons?

Summary of Working Paper No. 501 William B Francis

While the financial crisis had an adverse effect on the UK banking sector overall, some institutions fared worse than others in dealing with the onset of economic stresses. Those that fared worse were forced to undertake a host of more intensive actions, including debt-equity swaps (a form of bail-in), mergers with/acquisitions by stronger competitors and outright closure. But what was it about these firms that made them less capable of dealing with the downturn and what can regulators learn from these cases?

Toward addressing these questions, this paper takes a closer look at what drove UK deposit-takers' responses to the crisis. It specifically investigates the role that firm-level financial profiles played in influencing the intensity of such responses. It uses data spanning 2005 to 2011 on UK building societies, which, because of their mutual status, face similar constraints in their ability to tap external capital markets. This approach can help isolate the effect of financial condition, as opposed to market access, on response intensity.⁽¹⁾

The study groups firms into two separate and distinct categories according to the intensity with which they responded to the crisis. The first includes firms that resorted to more intensive efforts (ie debt-equity swap, mergers, acquisition, closure), while the second is effectively a catch-all category, consisting of firms that responded in other, less intensive, ways. It uses well-known empirical techniques (ie limited dependent variables models) and financial attributes from the research examining the determinants of bank failure/distress to investigate whether these factors are also useful in explaining UK deposit-taker response intensity. The financial factors examined include the well-known CAMEL attributes that analysts typically use to evaluate the condition of deposit-takers and that previous research finds useful in profiling banking institutions: Capital adequacy, Asset quality, Management capability, Earnings performance and Liquidity.

The paper's key result is that a small set of these financial attributes effectively distinguishes firms that undertook less intensive responses (ie less vulnerable firms) from those that resorted to more intensive responses (ie more vulnerable firms) to deal with the onset of economic stress. I also find

that, compared with risk-based capital measures, a simple leverage (ie capital to assets) ratio was better at classifying response intensity and, therefore, characterising financial vulnerability under the prudential regulatory regime that existed before the crisis. This evidence supports the recent regulatory emphasis on updating the regime to include consideration of non risk-based capital measures alongside risk-based measures.

A useful aspect of the modelling approach discussed in this paper is its objective consideration of a broad set of financial attributes and their interactions in profiling firm-level vulnerability. This approach means, for example, that low capital ratios would not be the sole criterion for triggering heightened supervisory attention. Rather, concerns about an institution's ability to deal with stress would be based on the financial CAMEL attributes as a group and their relative importance in explaining how firms responded to previous economic downturns. The output from the approach could also complement regular stress-testing efforts and assist in evaluating firms' recovery plans by pointing to firms that exhibit features similar to those that were less capable of dealing with the onset of adverse economic conditions in the past.

While the profiling approach discussed in this paper may be of interest to regulators for use in off-site monitoring, a key caveat limits its use in that capacity. In particular, the estimates in this study are conditioned on a prudential regime that excluded a leverage requirement. This study's findings, as a result, reflect UK deposit-taker behaviour that could conceivably differ from that under a regime that includes such a requirement (eg Basel III). This means that the set of financial measures — and the relative importance of each measure — found useful in distinguishing relatively more vulnerable firms in this study may be different under a revised prudential framework if deposit-takers alter business models and capital management practices in response. Still, the results are useful for highlighting potential shortcomings of the pre-crisis regulatory regime and for gaining initial insight into the effects of proposals aimed at addressing such flaws.

(1) Extending this analysis to include data from the wider UK banking sector is an area for future work.

The effect of the financial crisis on TFP growth: a general equilibrium approach

Summary of Working Paper No. 502 Stephen Millard and Anamaria Nicolae

The aim of monetary policy is to keep inflation low and stable. A major influence on inflationary pressure is the balance between an economy's capacity to supply goods and services — potential output — and the demand for these goods and services. In the wake of the financial crisis output in the United Kingdom fell dramatically while labour productivity fell initially and remains about 5% below its pre-crisis peak. This paper aims to show how a financial crisis might have a permanent impact on supply, specifically looking at total factor productivity (TFP): the element of productivity that cannot be explained by increases in inputs, particularly capital.

We use a simple growth model in which the growth rate is not fixed, but determined within the model, specifically by research and development (R&D) spending and the innovation that results from this. In this model a financial shock leads to a rise in the spread between the rate of interest paid by firms and the risk-free rate. Since firms in the model have to borrow to finance their R&D spending, the rise in the spread leads to a fall in R&D spending, which affects innovation and, hence, reduces TFP growth. In turn, this leads to permanent falls in the levels of output and labour productivity.

The key question for this paper is, then, to what extent the model suggests that the financial crisis can account for the weakness in UK productivity since the crisis via this channel. We would not expect the model to account for all of the fall in productivity as it leaves out, for example, the potentially long-lasting effects on productivity of impediments to capital being reallocated from less productive to more productive uses, the temporary effects of labour hoarding over the recession and of a labour supply response to the recession, the direct contribution of the financial sector to UK productivity, and the contribution of the oil and gas extraction sector (ie North Sea Oil), whose productivity was falling since before the crisis began. In addition, the effects in the model are likely to happen too quickly relative to the real world given that the lags between spending on R&D and the innovations resulting from such spending are likely to be much longer than the one quarter assumed in the model.

To be more specific, we perform the following simple experiment. We first construct a series for a 'financial shock' that replicates what happened in the United Kingdom in the wake of the financial crisis. We then run that shock process through the model and examine the implications for the endogenous variables of the model: in particular, labour productivity and TFP. We then compare these outcomes with the UK data on labour productivity.

The model suggests that we might expect the financial shock to lead to falls in GDP, TFP and labour productivity and that we would have expected several quarters of negative labour productivity growth, as we saw in the United Kingdom. However, the model fails to match the quantitative response of labour productivity growth suggesting a fall in average quarterly productivity growth of less than 0.05 percentage points during this period as compared with a fall in average productivity growth of just over 0.5 percentage points in the UK data.

We suggest several reasons why the modelled productivity response to the financial shock operating through this channel is quantitatively so small. First, it is not clear that we have managed to capture the full impact of the financial crisis on bank lending as it is likely that we saw an increase in quantitative constraints on borrowing, over and above the rise in spreads that drives the results. Second, the response of innovation to a given fall in R&D spending is likely to be much larger in the data than it is in our model on account of the fact that the general increase in uncertainty about demand that has been apparent since the crisis, and that is likely to act as a disincentive to innovation, is simply not modelled. If we put through our model a fall in innovation similar to that seen in the UK data, we are able to explain roughly 15% of the lower-than-expected UK labour productivity growth since the financial crisis. Adding in the effects of the financial shock on consumption and investment would probably help explain more of the short-run fall in productivity, as would allowing for an effect coming through working capital costs.

Peering into the mist: social learning over an opaque observation network

Summary of Working Paper No. 503 John Barrdear

Communication pervades human existence, and economic behaviour is no exception to this rule. In addition to the myriad of cultural interactions, people directly share economic information such as job opportunities and prices, and indirectly reveal information to each other as they trade goods and services. The study of how information is shared over a network of interactions is therefore an important field of economic research.

The topic of social learning — examining if, how, and how quickly people's beliefs might converge — when people communicate via a network has been examined extensively in the microeconomic literature. There has been little to no application to questions of macroeconomics, however, despite the common acceptance that imperfect access to information is critical to explaining the movement of aggregate variables. For example, firms' price-setting decisions may be influenced by observing each other's individual prices.

The reason that network learning has not been previously explored within macroeconomics is that three other features commonly deemed essential to the discussion of macroeconomics — that agents act repeatedly; that agents act strategically, with their pay-offs a function of other agents' actions; and that although imperfectly informed, agents' expectations are (close to) rational — make comprehensive analysis of network learning intractable in anything other than trivially small networks.

This paper presents a solution to this problem by proposing a simplifying assumption: that the network is 'opaque' in that economic players ('agents') such as households or firms do not know exactly who is connected to whom. Instead, it is supposed that agents know the probability distribution from which everybody draws the identity of their observees. That is, it is known that agent 1 is observed with a specific probability, agent 2 is observed with another probability, etc. The model also includes a key feature of actual networks by supposing that while most agents are unlikely to be observed, some groups of agents are disproportionately highly observed, even as the number of people in the network becomes very large.

Agents are attempting to learn about an unobserved or hidden 'state' variable (eg, the level of demand) by observing each other's actions. In the paper, the way that the possible expectations of this state (the 'hierarchy' of expectations) adjust over time is derived. With an opaque network, the hierarchy includes the average expectation regarding the hidden state, the average expectation of the average expectation, etc, but also includes an infinite sequence of weighted-average expectations and higher-order combinations between them.

Following a shock to the hidden state, average expectations respond more quickly than they do when agents do not observe each other in a network, but also temporarily overshoot the truth in a kind of herding behaviour that relies on the agents' observations of each other and their strategic motives (strategic meaning that they act taking into account beliefs about how others will respond). The degree of persistence of expectations is shown to be increasing in the number of competitors observed.

Idiosyncratic shocks (that is, those that affect only individual agents), which in many models have no effect on aggregate variables, are shown to influence the hierarchy of aggregate beliefs. Even when idiosyncratic shocks last only one period, these effects are also shown to be persistent, lasting for several periods. The paper therefore contributes to a new field of research demonstrating that aggregate volatility may emerge from idiosyncratic shocks.

Because of the focus on a setting with an underlying state that evolves over time and the way the hierarchy of average expectations evolves, those interested in exploring models of this type are able to determine the aggregate effects of network learning without a need to simulate individual agents' decisions. This makes the model particularly amenable to nesting within broader general equilibrium models of the economy that take account of all the interactions within and between different sectors of the economy — in other words, macroeconomic models.

Quantitative easing and bank lending: a panel data approach

Summary of Working Paper No. 504 Michael A S Joyce and Marco Spaltro

In response to the sharp deterioration in the global financial crisis in Autumn 2008, the major central banks cut their policy rates dramatically and began looking for other unconventional measures to loosen monetary conditions further. In the United Kingdom and United States, a key element of these unconventional measures has been the policy of large-scale asset purchases financed by central bank money, sometimes referred to as quantitative easing (QE).

In the United Kingdom, the Bank of England's Monetary Policy Committee (MPC) announced the introduction of the QE policy in March 2009, at the same time as it reduced Bank Rate to 0.5%, a historical low. In announcing the new policy, the Committee noted that without further measures there was a serious risk inflation would undershoot the 2% consumer prices index inflation target in the medium term. By the end of the first round of purchases that ended in January 2010 the Bank of England had purchased £200 billion of assets, consisting almost exclusively of government bonds — an amount equivalent to 14% of annual nominal GDP. In October 2011, the Bank resumed its QE purchases and by November 2012 the Bank had completed a further £175 billion of purchases.

There is now a large and growing literature that attempts to measure the impact of central bank asset purchases during the financial crisis in the United Kingdom and elsewhere. So far, the vast majority of research on QE has focused on its impact on economic growth and financial markets, while the effect of QE on bank lending has received much less attention. This relative neglect reflects the fact that policymakers in the United Kingdom and elsewhere expected QE to affect demand mainly through its impact on asset prices, while the effect on bank lending was expected to be small because of banks' incentives to deleverage and reduce the overall size of their balance sheets. This reasoning is consistent with the literature on the so-called bank capital channel, which suggests that capital can be an important driver of banks' lending decisions particularly in periods of market stress.

The MPC's caution about the strength of the bank lending channel was reflected in the design of the Bank of England's asset purchase programme, which was targeted towards the non-bank financial sector by skewing purchases towards medium and long-term maturity government securities (gilts), rather than the shorter-maturity gilts typically held by banks for their liquidity needs. However, to the extent that the Bank's QE asset purchases came from non-banks (directly or indirectly), the banking sector will have gained both additional reserves and a corresponding increase in its deposits. The additional reserves mean that banks' holdings of liquid assets will have increased, which might make banks more

willing to extend illiquid loans. At the same time, by increasing their deposits, QE will have made banks less reliant on seeking other funding to manage their liquidity needs. Put another way, the extra deposits that banks consequently held will have helped relieve any funding constraints they may have faced. Since these constraints are more likely to bind in times of financial stress, it seems possible that this might have led to additional lending. While any effects on lending might have been expected to be weak during a period when the banks were also trying to deleverage, it seems unlikely that there will have been no effect at all. In other words, relative to the counterfactual of no QE, bank lending seems likely to have been larger.

The contribution of this paper is to test for the existence of this bank lending channel historically and thereby to quantify the likely size of the effects of the Bank of England's QE policy during 2009–10 on bank lending, using a new non-publicly available quarterly panel data set on UK banks. The use of this unique data set allows us to model the relationship between bank lending growth and its determinants over a 20-year period pre-dating the financial crisis and to explore whether the relationship between deposits and bank lending changed during the crisis. We are also able to explore heterogeneities between large and small banks and to control for balance sheet effects, by including information on bank capital ratios at the level of individual banks. Using the historical relationships between bank lending growth and deposit growth, macroeconomic indicators and individual controls, we can then simulate the potential effects of QE on the banking sector.

We find that historically movements in the deposit ratio have a small but statistically significant effect on bank lending growth, which suggests that QE may have led to an increase in bank lending through its effect on deposits. These effects, however, are likely to have been small, both because the estimated marginal effects through deposits are small and also because we assume as a benchmark that there was a full pass-through from QE to deposits, which seems likely to overstate the impact. We also find no evidence that the impact from deposits increased during the QE period. Our analysis suggests that the effects on bank lending were heterogeneous across banks, as we find lending by small banks to be more responsive to the level of deposits than the lending of large banks. We also find evidence that bank lending is positively related to how well capitalised banks are, suggesting that the impact of QE on bank lending may have been weaker because of the lower levels of capital during the crisis. In a sense, this is to be expected and justifies the emphasis policymakers gave to QE going round the banks. At the same time, it suggests that macroprudential policy may potentially influence the effectiveness of monetary policy.

The cost of human capital depreciation during unemployment

Summary of Working Paper No. 505 Lien Laureys

Unemployment is an important driver of potential supply, and of crucial interest to policymakers for this reason, as well as the effect on the well-being of households. It is well understood that one undesirable aspect is that skills — human capital — may deteriorate as unemployment spells lengthen. This paper analyses how this human capital depreciation affects the efficiency of aggregate labour market outcomes. This may help us to understand the dynamics of unemployment better.

The framework of analysis is an otherwise standard model of search for jobs by the unemployed to which human capital depreciation is introduced. Workers who had their human capital eroded while being unemployed are less productive upon re-employment than workers who were not so affected. At the same time, it allows for learning-by-doing such that workers with depreciated human capital can regain skills while being employed.

In the presence of human capital depreciation during unemployment, firms' hiring decisions affect not only the unemployment rate, but also the share of workers with eroded skills in the unemployment pool. Hiring therefore influences workers' chances of finding jobs, average unemployment duration, and thus the extent of skill erosion. For example, when firms hire less, unemployed workers have a smaller chance of finding a job, which increases their unemployment duration. Longer unemployment spells in turn raise the probability that their human capital erodes. As a result, a drop in hiring increases the relative share of job-seekers with eroded skills in the unemployment pool.

In the model, it is assumed that the unemployment pool's skill composition determines how likely it is that job-seekers with or without eroded skills show up for job interviews. Thus, the pool's composition determines the average productivity of job candidates. Consequently, firms' hiring decisions, through their effect on job-seekers' skills, affect the output that can be

generated by other firms' new matches. This amounts to a composition externality (a cost or benefit imposed on other firms) related to job creation, which arises in addition to the familiar congestion externality following from the search frictions (whereby an extra unemployed person makes it harder for other unemployed workers simply because there are more people searching). The composition externality arises because firms ignore how their hiring decisions today affect the unemployment pool's skill composition in the next period, and hence the expected productivity of other firms' new hires. As a result, when human capital depreciates during unemployment, there are gains from job creation which are not fully internalised.

Insight into the composition externality may be provided by analysing the policy instrument that can replicate a hypothetical planner's solution when this externality is the only source of inefficiency, and financing goes through non-distortionary taxation. In the model, the instrument takes the form of a state-dependent employment subsidy implying that because of this externality job creation in the *laissez-faire* economy is too low in all states of the economy from a social point of view. But the extent to which job creation is too low varies over the cycle. This is because the externality's magnitude, which depends on the impact of job creation on the pool's skill composition, reduces when the share of unemployed workers who already have eroded skills increases. How this externality's magnitude varies over the cycle depends on the dynamic path of human capital depreciation, as this will influence the point in the cycle at which this share starts to increase.

Calibrating the model to the US economy shows that the composition externality is quantitatively relevant. When skill loss is the only source of inefficiency, restoring constrained-efficiency entails a drop in the average unemployment rate in the range of 0.92 to 0.27 percentage points.

Tailwinds from the East: how has the rising share of imports from emerging markets affected import prices?

Summary of Working Paper No. 506 John Lewis and Jumana Saleheen

Over the past fifteen years, there has been a significant rise in the share of UK imports coming from industrialising or emerging market economies (EMEs), such as China, India and the new EU member states of Central and Eastern Europe. Since these countries typically have much lower prices and wages, policymakers and academics have argued that the growing share of imports from EMEs has pushed down on import prices in developed economies.

Our goal is to quantify the impact of the rising share of EME imports on import prices in the United Kingdom. The argument runs that as 'cheap imports' displace the products of industrialised countries with cheaper goods from EMEs it will push down on aggregate import prices. This happens partly as importers 'switch' to the cheaper goods from lower-wage economies — we term this the 'switching effect'. And partly because producers from other countries lower their prices in response to the increased competition from EMEs — we term this the 'competition effect'. But there is another potentially countervailing effect that has gained attention recently. It relates to the observation that EME inflation has been higher than developed economies recently, so greater exposure to EMEs would lead to upward pressure on import prices — we call this the 'inflation effect'. This paper investigates the size of each of these channels.

A rising share of imports from EMEs may also feed through to affect a broader set of producer and consumer prices; either because of competition effects or because imports are used as in intermediate input in the production process. In this paper we focus only on import prices.

We think that impact comes through the three main channels described above, and we seek to quantify the size of each.

Our main data source is the UK customs authority (HMRC), which includes data on both the volumes and values of imports, by country of origin for over 3,000 industries, 2,000 of which are in manufacturing. This highly disaggregated data allows us to account for differences across industry groups. We also allow for different effects across EME country groups by dividing our sample of EMEs into China, the new EU member states and other low wage cost economies such as India.

We find that when China gains market share in an industry, import prices do tend to fall, although this effect differs across industries. For the other EME country groups we find no clear link between gaining market share and lower import price inflation. We also find little evidence for the 'inflation effect'. Overall, that implies that emerging economies have lowered, rather than raised import price inflation in the United Kingdom.

The finding that China exerts that largest and statistically significant downward impact on UK import prices reflects the fact that China has gained market share more quickly than other EMEs and, that China has a lower price level than most other EMEs. We estimate this 'tailwind' from China has lowered UK import price inflation by around 0.5 percentage points per year. Although there is some variation from year to year, there is no discernible trend over time, so we conclude that the tailwinds from China were blowing just as strongly in 2011 as they were a decade ago.

Estimating time-varying DSGE models using minimum distance methods

Summary of Working Paper No. 507 Liudas Giraitis, George Kapetanios, Konstantinos Theodoridis and Tony Yates

Much modern macroeconomic research and policy analysis is predicated on the idea that the model is 'stable over time'. What we mean by this is that the structural parameters (ie, 'deep' determinants such as households and firms' preferences, the nature of production functions, how prices are set and properties of the random shocks that constantly buffet the economy) are constant over time. Models are estimated invoking this assumption and then used to explain past macroeconomic data or to forecast the future.

However, this assumption of 'constancy' is just that: an assumption. A literature has grown up that looks into this parameter constancy, and often finds that empirically it appears not to hold. This paper contributes to this effort. A standard empirical time-series model is estimated on US data where every variable in the system is a function of all lagged variables in the system (known as a vector autoregressive model) but where the theory-free non-structural parameters of this empirical model are allowed to vary with time. The next step is to estimate a popular theoretical model, spelling out the economic theory with a specific structural parameterisation used by many academic researchers and central banks by choosing its parameters so the theoretical model displays dynamic responses to shocks that match those predicted by the empirical model as closely as possible. This is done for every period in the sample, as the time-varying parameters of the time-series model define responses that are different for every period in the sample.

It emerges that there is substantial variation in key parts of the model. These include the 'stickiness' that determines the speed of adjustment of prices and wages; the speed with which investment responds to changes in the user cost of

capital; and changes in the determinants of how swiftly consumption responds to shocks.

These parameters have been the focus of criticism before, from economists that associate themselves with the view that macroeconomies are relatively frictionless, and argue they lack independent empirical evidence that justify their existence in the theoretical model. So the fact that they move around a lot over time might be taken as evidence to reinforce their scepticism. Furthermore, models that change markedly over time could simply be misspecified. In which case, our results suggest, echoing findings from previous papers, that there is work to do to dig deeper in those aspects of the macroeconomy that give rise to this apparent time variation in the parameters.

On the other hand, if one is prepared to accept the notion of time-varying theoretical models, they can be put to work to see whether they change the answers to questions that were previously only posed in the context of fixed-parameter models. For example, the parameters that define monetary policy behaviour moved less than has previously been suggested. There is no dramatic difference in the estimates between pre and post-Volcker monetary policy; the dramatic difference in performance is explained as a difference between the variance of supply shocks over the two periods. As another example, there are substantial fluctuations in the contributions of different shocks at different time periods to the business cycle. This might explain some of the controversy in the fixed-coefficient literature that has looked at the same issue, using different data sets and different time periods. So all this suggests that time variation has important implications for policy.

How does credit supply respond to monetary policy and bank minimum capital requirements?

Summary of Working Paper No. 508 Shekhar Aiyar, Charles W Calomiris and Tomasz Wieladek

Traditionally national authorities have regulated banks from the perspective of the safety and soundness of individual institutions. Such 'microprudential' regulation has operated separately from the main policy instrument employed to smooth aggregate fluctuations in business activity, monetary policy. But following the recent global financial crisis, 'macroprudential' regulation, such as varying banks' capital requirements countercyclically, has increasingly been viewed as a desirable instrument of policy. Changing banks' capital requirements countercyclically not only has the familiar aim of building up capital in good times to act as a buffer to absorb losses in bad times, it also can have the goal of stabilising the credit cycle itself, leaning against the cycle to reduce credit growth when the economy overheats, and mitigating disruptive credit crunches when the economy suffers a downturn. This latter goal is appropriately 'macroprudential', since a shallower credit cycle should reduce the incidence of financial crises generated by imprudent lending and the mispricing of risk, thus enhancing the stability of the financial system. But higher capital requirements could also increase lending at banks with very low or negative net worth, in particular if they helped to overcome a so-called 'debt overhang' problem.

There is already a substantial and rapidly growing theoretical literature on the expected credit supply impact of bank capital requirements (alongside the venerable literature on the credit supply impact of monetary policy). Moreover, some papers predict that monetary policy should interact with changes in bank capital requirements through various channels when the two instruments are deployed jointly. That is to say, a bank's lending response to a change in capital requirements may be different if there is a simultaneous change in monetary policy, and a bank's lending response to a change in monetary policy may be different if there is a simultaneous change in capital requirements. So far, however, there have been no empirical

tests of whether or not this is the case, despite their evident and urgent relevance to policy.

This paper provides the first empirical estimate of how banks' credit supply responds to monetary policy and minimum capital requirements, when the two instruments are used together. The analysis is made possible by an apparently unique policy experiment performed in the United Kingdom during the 1990s and 2000s, where the Financial Services Authority varied individual banks' minimum risk-based capital requirements. The extent of this variation across banks was large (the minimum required capital ratio was 8%, its standard deviation was 2.2%, and its maximum was 23% of risk-weighted assets). The variation in the average minimum capital requirement over the business cycle was also large, and tended to be countercyclical, as envisaged under macroprudential regulation. This data set on individual banks' minimum capital requirements over time is combined with Bank of England data on lending by the same banks.

The empirical analysis suggests that tightening monetary policy and increasing banks' minimum capital requirements both have independent negative effects on banks' supply of loans to the non-financial private sector. Consistent with previous work it is found that lending by large banks does not react as much as the lending of small banks to changes in monetary policy, perhaps because large banks have greater flexibility in accessing non-deposit funding. Changes in capital requirements, on the other hand, have large effects on the loan supply of large and small banks alike, suggesting greater relative potency for this instrument in economies with banking systems comprised of a small number of large banks. Finally, contrary to existing theoretical perspectives on the interaction of monetary policy and capital requirement changes, no interaction effects are found between changes in monetary policy and capital requirements.