BANK OF ENGLAND

Working Paper No. 502

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

Stephen Millard and Anamaria Nicolae

June 2014

BANK OF ENGLAND

Working Paper No. 502
The effect of the financial crisis on TFP growth: a general equilibrium approach

Stephen Millard[(1)] and Anamaria Nicolae[(2)]

Abstract

In this paper, we use a simple endogenous growth model to show how a financial crisis might have a permanent effect on the level of total factor productivity (TFP). In the 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 have to borrow to finance their research and development (R&D) spending, such a 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.

Key words: Endogenous growth, research and development, innovation.

JEL classification: O4.

(1) Bank of England, Durham University Business School and Centre for Macroeconomics. Email: stephen.millard@bankofengland.co.uk

(2) Durham University Business School. Email: anamaria.nicolae@durham.ac.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. The authors wish to thank Nick Oulton, Peter Sinclair, seminar participants at the Bank of England, University of Hull and Durham University Business School and participants at the ESRC Money, Macro and Finance Research Group Annual Conference held in Dublin, September 2012, and the Growth and Policy Conference held in Durham, March 2013, for helpful comments. Any errors and omissions remain the fault of the authors. This paper was finalised on 5 June 2014.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Contents

Summary

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 re-allocated 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 outturns 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.

1 Introduction and motivation

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 inflation remained above its target, as shown in Chart 1. The behaviour of labour productivity (measured as GDP divided by the Labour Force Survey measure of employment) during and after the crisis, shown in Chart 2, suggests that the economy has experienced a hit to supply. The aims of this paper are to explain one channel through which a financial crisis might have a permanent impact on supply, specifically looking at Total Factor Productivity (TFP).

Chart 1: Annual GDP growth and CPI inflation

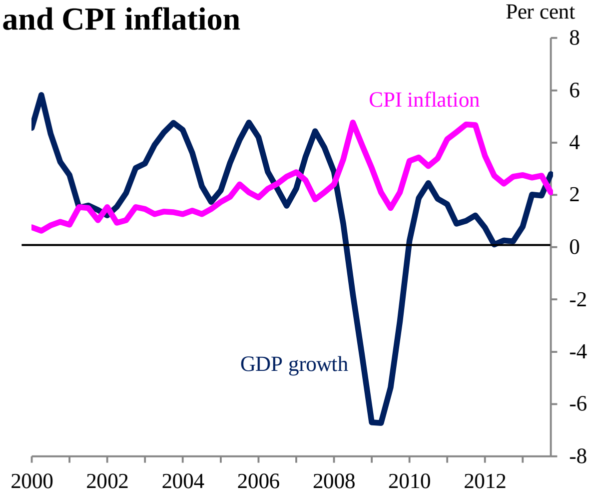

Chart 2: Labour productivity

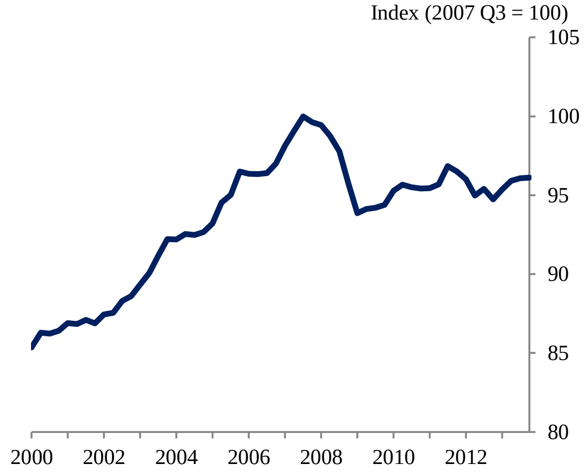

Oulton and Sebastiá-Barriel (2013) have carried out some empirical work that shows that financial crises can permanently affect the level of labour productivity. Specifically, they used data on productivity derived from the Conference Board's Total Economy Database for 2011 together with the Reinhart and Rogoff (2009) data on financial crises: in total 61 countries covering the period 1950-2010. They found that a financial crisis temporarily reduces productivity growth by about one percentage point, with a permanent effect on the level of productivity. They suggest that this happens because financial crises impede the reallocation of capital from less productive to more productive uses. Barnett *et al.* (2014 a and b) use firm-level data to look more closely into this channel and find evidence that capital reallocation has, indeed, been impaired since 2008 and is likely to have contributed to the weakeness in UK productivity over this period.

The current paper, however, examines a different, though complementary, channel through which a financial crisis can permanently affect the level of TFP. More specifically, we use a simple endogenous growth model in which 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 this model have to borrow to finance their research and development (R&D) spending, such a rise in the spread leads to a fall in R&D spending; this affects innovation and, hence, reduces TFP growth. In turn, this leads to permanent falls in the levels of output and labour productivity.

This channel seems to be at play in the UK data. Chart 3 plots real business R&D spending. In response to the recessions of 1989/90 and 2008/9, real R&D spending fell. And both of these recessions are characterised as financial crises by Oulton and Sebastiá-Barriel (2013) suggesting a possible link between financial crises and R&D spending. Furthermore, we can look at some survey evidence on innovation: specifically, the *UK innovation survey* conducted by the Department for Business, Innovation and Skills in the United Kingdom. Chart 4 suggests that both product and process innovation fell between 2009 and 2011, suggesting a link between financial crises and innovation.

Chart 3: Real business spending on research and development

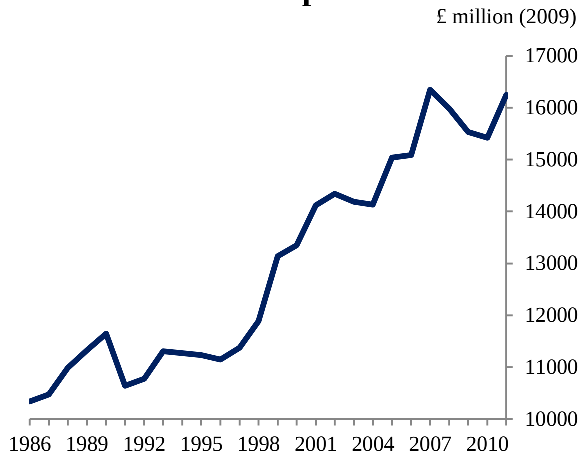

Chart 4: Innovation

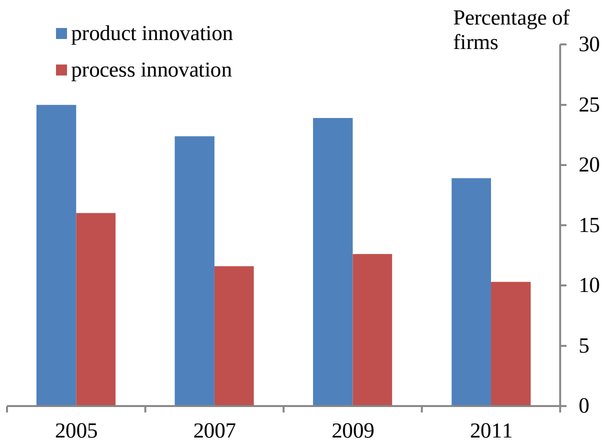

Given our model, we can ask ourselves to what extent it suggests that the financial crisis can account for the weakness in UK productivity we have seen since the crisis via this channel. We would not expect it to account for all of the fall in productivity as it leaves out, for example, the temporary 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 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. Nevertheless, the exercise is worthwhile as a guide as to the extent to which the financial crisis might have reduced innovation and permanently damaged the level of output and productivity, an important question to be able to answer in the monetary policy context.

The structure of the paper is as follows. In the next section, we review the literature on financial frictions, growth and productivity. Section 3 lays out our model while Section 4 discusses how we use the steady-state implications, together with UK data, to calibrate the key parameters. Section 5 presents some implications for stylised facts governing R&D, innovation and productivity and Section 6 performs an experiment in which we analyse the effects of the financial crisis as proxied by movements in spreads in the United Kingdom since 2008. Section 7 concludes.

2 Previous literature

Our work brings together two strands of literature. One strand looks at how a shock to credit can affect the supply of output (as opposed to demand) and the other looks at endogenous growth models, in which temporary reductions in, say, investment can permanently affect output. The familiar credit channel models of Bernanke and Gertler (1989), Kiyotaki and Moore (1997) and Bernanke *et al.* (1999) are about how issues around credit availability can affect demand in the economy and the way the economy responds to shocks; aside from the obvious effect on the capital stock coming through investment, these are not really models of how credit conditions affect supply. It is this alternative strand of literature that is more relevant to this paper.

2.1 Channels through which a shock to credit affects supply

Credit can affect the ability of firms to supply output through (at least) three channels: the cost and availability of working capital; the cost of intangible investment (which may in turn affect innovation); and the degree of firm entry and exit.

The working capital channel was examined by, among others, Fernandez-Corugedo *et al.* (2011). Intuitively, we can think of credit used for working capital purposes as an additional input in the production function; if working capital becomes more expensive, then firms will have to reduce their output, raise their prices or both. The credit contraction associated with the financial crisis had a significant effect on the working capital position of firms, and this is likely to have affected potential supply, at least temporarily. Fernandez-Corugedo *et al.* (2011) found that disruptions to the supply of credit have had large and persistent effects on the real economy, through the working capital channel, though in their model, credit disruptions cannot have a permanent effect on the level of output.

Credit frictions can also affect the cost of 'intangible' inputs and, so, measured TFP. Dal Borgo *et al.* (2013) document intangible investment for the United Kingdom and see how it contributes to economic growth by employing the framework of Corrado *et al.* (2005, 2009) to a new set of data based on a micro survey. They found that, for 2000-2008, intangible capital deepening accounted for 23% of productivity growth, against computer hardware 12% and TFP 40%. Once intangibles are included in the growth accounting framework, estimated TFP growth is lowered by 15%. The relevance for us is that any change in the cost and availability of credit, particularly unsecured credit (since intangible capital cannot be offered as collateral) is then likely to have an impact on productivity growth.

Finally, credit frictions can also affect supply via their effect on the entry and exit of firms. Disney *et al.* (2003), measure and examine the contribution of internal and external restructuring to productivity growth. Internal restructuring consists of introducing new technology and organisational change, and external restructuring consists of entry, exit, and market share change. The idea is that increasing internal restructuring to allow the use of new technology, and the replacement of exiting low productivity firms by higher productivity entrants and higher productivity firms gaining market share, leads to productivity growth. They show that external restructuring accounts for 50% of establishment labour productivity growth and 80-90% of establishment TFP growth. Given these results, it is clear that credit frictions can affect productivity growth via their effect on firm entry and exit. In particular, if low productivity firms survive for too long as a result of bank forbearance, and potential entering firms are unable to gain the credit they need to set up, then productivity will be lower than if the credit market was functioning properly. This channel, first identified in Japan, is increasingly gaining in attention. Caballero *et al*. (2006), for example, show how the existence of zombie firms in Japan interfered with the process of creative destruction and stifled growth by distorting competition throughout the rest of the economy. The zombie firms contributed to depressed market prices for their products, raising market wages, and congesting the markets. At the same time, new firm entry was prevented, especially in the sectors affected by the large stock and land price declines like construction and real estate. Along similar lines, Kiyota *et al.* (2005) found that between 1994 and 1998, efficient firms in Japan were exiting while inefficient firms survived leading to a substantial fall in TFP after 1996. This suggests that 'natural selection' does not function during a financial crisis induced recession.

2.2 Models of endogenous growth

Turning to endogenous growth models, Aghion and Howitt (2000) discuss three main branches. One is the R&D-based growth model of Schumpeter (1942), one is based on increases in product variety, and the other is the capital-driven growth model. Within the R&D-based growth models, Comin and Gertler (2003), Barlevy (2007) and Nuño (2011) consider models that embed endogenous technological progress, which is one determinant of productivity, into standard DSGE models. Also within this branch of the literature, Lee and Mukoyama (2008) and Jaimovich and Floetotto (2008) examine entry and exit of firms over the business cycle and how this links to mark-ups and productivity.

But for the current paper, we are most concerned with those endogenous growth models that attempt to explain the links between long-run economic growth and financial intermediation. In one such, Michalopoulos *et al.* (2009) show that unless financiers continually innovate, economic growth stops. They model technological and financial innovation as reflecting the decisions of profit maximizing agents and explore their implications for economic growth.

Dekle and Kletzer (2003) adopt a general equilibrium endogenous growth model with financial intermediation to show how government policies towards the financial sector can lead to banking crises and persistent growth slumps. Dekle and Kletzer (2004) take a similar approach with bank-intermediated investment, showing how public deposit insurance and weak prudential regulation (government failure to reinforce the writing off of non-performing assets) combined can affect investment and output growth, leading to banking crises and permanent declines in growth. The assumptions of both theoretical models are based on essential features of the Japanese financial system and its regulation. Although these papers bring forbearance into an endogenous growth framework, their focus is on government failure to intervene to address the problem of misplaced loans rather than on the effect of reduced credit availability on long-run growth.

3 The model

In this section, we present a theoretical framework to explain the effects of a financial crisis on productivity, with the aim of assessing to what extent the recent crisis can explain the recent fall in UK productivity. As we are interested in long-run productivity rather than medium-run movements in output, inflation and productivity, we develop an entirely 'real' model in which there is no role for monetary policy. This greatly simplifies the analysis without affecting the channels in which we are interested. We present in turn the problems faced by the different agents we consider: consumers, final goods producers, intermediate goods producers and banks. The model is taken directly from the 'endogenous growth' literature linking financial constraints with R&D and productivity growth and closely follows that laid out in Aghion and Howitt (2000), Chapter 6, Section 2.

3.1 Consumers

Consumers solve a standard maximisation problem. They maximise utility, which is increasing in consumption, c_t, and decreasing in hours worked, h_t, subject to a budget constraint:

$$\text{Max} \quad \sum_{t=0}^{\infty} \beta^t \left(\ln c_t - \frac{\vartheta}{1+\sigma} h_t^{1+\sigma} \right)$$

$$\text{Subject to} \quad D_t = (1+i_{t-1}) D_{t-1} + w_t h_t + \Pi_t - c_t \qquad \textbf{(1)}$$

where β is the discount factor, ϑ scales the marginal disutility of work, σ is the inverse Frisch elasticity of labour supply, D denotes real, risk-free, bank deposits at the end of time t, i is the risk-free real interest rate, w is the real wage, and Π denotes the real profits of the firms and the banks, which are distributed lump-sum to the consumers who are assumed to own them.

The first-order conditions imply:

$$1 = \beta (1+i_t) E_t \left(\frac{c_t}{c_{t+1}} \right) \qquad \textbf{(2)}$$

$$\vartheta h_t^{\sigma} c_t = w_t \qquad \textbf{(3)}$$

Equation **(2)** is the IS curve linking expected consumption growth to the real interest rate and equation **(3)** is the labour supply curve.

3.2 Final goods producers

Following Aghion and Howitt (2000), we suppose that final output is produced under perfect competition by combining labour and j differentiated intermediate goods, x_j.

$$y_t = h_t^{1-\alpha} \int_0^1 A_{jt}^{1-\alpha} x_{jt}^{\alpha} dj \tag{4}$$

where A_{jt} is TFP for the intermediate producer in sector j at time t, which will depend on whether or not it innovates new technology over the period, y_t is output of the final good in period t, x_{jt} is the amount of intermediate product from sector j used at time t and $\alpha \in (0,1)$. Hence, the firms' profit will be given by:

$$h_t^{1-\alpha} \int_0^1 A_{j,t}^{1-\alpha} x_{j,t}^{\alpha} dj - w_t h_t - \int_0^1 p_{j,t} x_{j,t} dj \tag{5}$$

where p_j is the relative (to the price of the final good) price of intermediate good j.

Maximising profit implies the demand for intermediate good j will be given by:

$$p_{j,t} = \alpha h_t^{1-\alpha} A_{j,t}^{1-\alpha} x_{j,t}^{\alpha-1} \tag{6}$$

The labour demand curve will be given by:

$$(1-\alpha) y_t / h_t = w_t \tag{7}$$

3.3 Intermediate producers

The key element of this model is that TFP in the intermediate sector is determined by the extent to which firms in that sector innovate. In turn, innovation will be determined by the extent to which these firms invest in R&D. Here we are using R&D as a proxy for investment in all sorts of intangibles plus experimentation in general including in new businesses processes and models, start-ups, job switching, etc. Consequently, we should note that measured R&D spending may not be picking up all the investment that improves firms' chances of innovating and may be picking up investment that has no effect on firms' chances of innovating. The key is that, in the model, this investment results in a greater likelihood of these firms innovating.

We first consider TFP for the intermediate producer in sector i. If this firm successfully innovates in period t, it is able to use the frontier technology, in which case its total factor productivity will be given by:

$$A_{i,t} = \overline{A}_t = (1+\gamma) A_{t-1} = (1+\gamma) \int_0^1 A_{i,t-1} di \tag{8}$$

where $\gamma > 0$ is the size of the innovation, ie, the rate at which the frontier grows relative to last period's average productivity. If an intermediate producer does not successfully innovate, its productivity will remain unchanged from the previous period: ie, will equal $A_{i,t-1}$.

So, by integrating over all sectors, we can show that average TFP at time t will be given by:

$$\begin{aligned} A_t &= \mu_t (1+\gamma) A_{t-1} + (1-\mu_t) \int_0^1 A_{i,t-1} di \\ &= (1+\gamma\mu_t) A_{t-1} \end{aligned} \tag{9}$$

where μ_t is the probability that a firm successfully innovates at time t.

The growth rate of average TFP will then equal:

$$g_t = \frac{A_t - A_{t-1}}{A_{t-1}} = \mu_t \gamma \tag{10}$$

The probability that an intermediate producer will successfully innovate in period t will depend upon its investment in R&D in period t-1. The more investment in R&D they carry out, the greater the probability of innovating during the following period. However, the greater is their current TFP, then the more R&D spending they would need to do in order to achieve a given probability of innovating during the following period. Following Aghion and Howitt (2000) we adopt the following functional form:

$$\mu_{j,t+1} = \rho \left(\frac{R_{j,t}}{(1+\gamma) A_t} \right)^{\phi} \tag{11}$$

where R is investment in R&D and ρ is a parameter that reflects the productivity of the research sector. We differ from Aghion and Howitt (2000) who allowed R&D spending to affect the probability of innovating during the current period. We can note that either of these approaches is likely to result in too short a lag between movements in the cost of R&D borrowing and movements in productivity growth, given the lags involved in investment in R&D and other intangibles coming to fruition. In a sense, though, this is no different to what is typically assumed about investment in physical capital.

The problem for an intermediate producer is to maximise the present discounted (utility) value of its cash flow to consumers subject to demand for its goods coming from the final goods producers, a simple production function linking output of the intermediate good to capital used in its production and the evolution of capital itself. Intermediate firms also spend money on R&D, which increases the probability with which they will be able to innovate during the following period, raising their productivity. We assume that intermediate firms have to borrow the money to do this.[1] At the beginning of the period they pay back their R&D loans from the previous period at an interest rate set during the previous period. During the period money flows into their accounts to pay for R&D spending and immediately flows out as the R&D budget is spent. So the actual flow of cash associated with paying for the R&D spending happens during the following period.

Putting all this together, we can write the problem for intermediate firm *j* as:

$$\text{Max} \quad E_0 \sum_{t=0}^{\infty} \frac{\beta^t}{c_t} \left(p_{j,t} x_{j,t} - I_{j,t} - R_{j,t-1} \left(1 + i_{L,t-1} \right) \right)$$

Subject to

$$p_{j,t} = \alpha h_t^{1-\alpha} A_{j,t}^{1-\alpha} x_{j,t}^{\alpha-1} \quad \textbf{(12)}$$

$$x_{j,t} = k_{j,t-1} \quad \textbf{(13)}$$

$$k_{j,t} = (1-\delta) k_{j,t-1} + I_{j,t} \quad \textbf{(14)}$$

And equation **(11)**, $\mu_{j,t+1} = \rho \left(\frac{R_{j,t}}{(1+\gamma) A_t} \right)^{\phi}$. Here *I* is investment in physical capital, i_L is the (real) loan rate, *k* is the end of period stock of capital and δ is its depreciation rate.

The firm will be choosing its end-of-period capital stock, R&D spending and, by implication, its probability of innovating. Noting this and substituting in the relevant constraints, we can rewrite the problem for the firm to choose k_t and R_t in order to maximise

$$E_0 \left(\begin{array}{l} \sum_{s=0}^{t} \frac{\beta^s}{c_s} \left(\alpha h_{j,s}^{1-\alpha} A_{j,s}^{1-\alpha} k_{j,s-1}^{\alpha} - k_{j,s} + (1-\delta) k_{j,s-1} - R_{j,s-1} \left(1 + i_{L,s-1} \right) \right) \\ + \frac{\beta^{t+1}}{c_t} \left(\begin{array}{l} \rho \left(\frac{R_{j,t}}{(1+\gamma) A_t} \right)^{\phi} \alpha h_{j,t+1}^{1-\alpha} \left((1+\gamma) A_{j,t} \right)^{1-\alpha} k_{j,t}^{\alpha} + \left(1 - \rho \left(\frac{R_{j,t}}{(1+\gamma) A_t} \right)^{\phi} \right) \alpha h_{j,t+1}^{1-\alpha} A_{j,t}^{1-\alpha} k_{j,t}^{\alpha} \\ - k_{j,t+1} + (1-\delta) k_{j,t} - R_{j,t} \left(1 + i_{L,t} \right) \end{array} \right) \\ + \sum_{s=t+2}^{\infty} \frac{\beta^s}{c_s} \left(\alpha h_{j,s}^{1-\alpha} A_{j,s}^{1-\alpha} k_{j,s-1}^{\alpha} - k_{j,s} + (1-\delta) k_{j,s-1} - R_{j,s-1} \left(1 + i_{L,s-1} \right) \right) \end{array} \right) \quad \textbf{(15)}$$

The first-order conditions for this problem, once we have integrated over all the intermediate producers, imply:

$$\alpha^2 E_t A_{t+1}^{1-\alpha} h_{t+1}^{1-\alpha} k_t^{\alpha-1} = i_t + \delta \quad \textbf{(16)}$$

[1] This reflects the stylised fact that most investment is financed out of retained earnings whereas most R&D spending is financed by bank borrowing.

$$\left(1+i_{L,t}\right)=E_t\alpha h_{t+1}^{1-\alpha}A_t^{1-\alpha}k_t^{\alpha}\left(\left(1+\gamma\right)^{1-\alpha}-1\right)\frac{\phi\mu_{t+1}}{R_t} \quad \textbf{(17)}$$

Equation **(15)** is the demand curve for capital. It sets the expected marginal revenue product of capital next period, $\alpha^2 E_t A_{t+1}^{1-\alpha}h_{t+1}^{1-\alpha}k_t^{\alpha-1}$, equal to its user cost, $i_t+\delta$, which is paid this period (since capital is put in place the period before it is actually used). Equation **(16)** sets the cost of investing in an additional unit of R&D, $1+i_{L,t}$, equal to the expected benefit. This amounts to the marginal increase in the probability of innovating next period TFP next period associated with the R&D spending, $\frac{\phi\mu_{t+1}}{R_t}$, multiplied by the marginal product of the associated innovation, $\alpha h_{t+1}^{1-\alpha}A_t^{1-\alpha}k_t^{\alpha}\left(\left(1+\gamma\right)^{1-\alpha}-1\right)$.

3.4 Banks

Banks set the loan rate in order to maximise profits. They fund their lending to intermediate goods firms for R&D, *R*, by accepting deposits from consumers, on which they pay interest at the risk-free rate. We assume perfect competition in banking, which implies the zero-profit condition:

$$\left(\theta_t+i_t\right)R_t=i_{L,t}R_t \quad \textbf{(18)}$$

Here θ_t is a monitoring cost paid by the banks. Unlike Aghion and Howitt (2000), we do not allow default to occur in equilibrium. Rather we assume that, as long as banks pay the monitoring cost, they can ensure that they only lend money to firms that will pay them back. One could think of the need to pay this monitoring cost arising from a model in which, out of equilibrium, there are a number of firms that, if allowed to operate, would, in fact, default; paying the monitoring cost enables banks to avoid lending to such firms. Alternatively, we can think of the monitoring cost as capturing, in a simple way, the efficiency of the banking system in being able to originate loans to firms that need them; the higher is the cost, the higher the return to the banks of making loans needs to be and, so, the fewer loans are viable.

Rearranging the zero-profit condition allows us to write the spread of loan rates over the risk-free rate as:

$$i_{L,t}-i_t=\theta_t \quad \textbf{(19)}$$

Given the earlier discussion, we should note that the spread is not really representing a 'risk premium'; rather it is capturing the 'costs of banking'. So, a financial shock, proxied by an increase in the cost of monitoring R&D loans (ie, the cost of banking), will lead to an increase in the spread between loan rates and the risk-free rate. If we think that the rise in spreads we have seen in recent years has been driven by a fall in the efficiency of the banking sector (ie, increase in the costs of banking), then we can use this equation to calibrate the size of the financial shock – ie, shock to the ability of the banks to intermediate – we have seen over recent years given data on the spread. A rise in the spread will lead to an increase in the cost of R&D, leading to a reduction in R&D spending and a permanent effect on total factor productivity. The effect on investment, on the other hand, will depend on what happens to the risk-free rate.

3.5 Market clearing and equilibrium

To complete the description of the model, we combine the various constraints faced by consumers and firms (together with the definitions of profits, etc,) in order to obtain the aggregate resource constraint:

$$y_t - \theta_t R_t = c_t + I_t + R_t \qquad \textbf{(20)}$$

That is, final output net of bank monitoring costs is equal to consumption plus investment in physical capital plus investment in R&D.

An equilibrium for this model is a set of allocations, interest rates, a wage and a relative price of intermediate goods such that consumers are maximising their utility subject to their budget constraint and firms and banks are all maximising their profits subject to their constraints. The equilibrium is described by equations **(2)**, **(3)**, **(4)**, **(6)**, **(7)**, **(9)**, **(12)**, **(13)**, **(14)**, **(16)**, **(17)**, **(18)**, and **(20)**. These equations solve for y, I, c, k, x, h, p, w, i, i_L, A, μ and R conditional on an assumed exogenous process for θ.

4 Steady state and calibration

In order to calibrate the parameters of our model, we solve for its steady state and then set the parameters so that some key steady-state ratios implied by the model match the average values of such ratios in the UK data.

4.1 Steady-state equations

We define a steady state as an equilibrium in which all real variables are growing at the rate g. For example, for TFP:

$$\frac{A_t - A_{t-1}}{A_{t-1}} = g, \forall t \qquad \textbf{(21)}$$

Equation **(9)** links this to the probability of innovation and the technological advantage gained by innovating:

$$g = \mu\gamma \qquad \textbf{(22)}$$

We use the IS curve linking expected consumption growth to the real interest rate given by equation **(2)** to obtain, for a given value of the parameter β, the risk-free real interest rate i:

$$i = \frac{1+g}{\beta} - 1 \qquad \textbf{(23)}$$

We normalise total hours worked to be unity in steady state. Then, the labour supply curve, equation **(3)**, implies:

$$\vartheta = \frac{w}{c} \tag{24}$$

Now, if we combine equations **(4)** and **(12)**, we obtain:

$$1 = \left(\frac{A}{y}\right)^{1-\alpha}\left(\frac{k}{(1+g)y}\right)^{\alpha} \tag{25}$$

Equation **(6)** implies the relative price of intermediate goods will be given by:

$$p = \alpha\left(\frac{A(1+g)}{k}\right)^{1-\alpha} \tag{26}$$

From the labour demand curve, equation **(7)**, we get:

$$(1-\alpha) = \frac{w}{y} \tag{27}$$

The capital accumulation equation implies:

$$\frac{I}{k} = 1 - \frac{1-\delta}{1+g} \tag{28}$$

From the capital demand curve, equation **(16)**, we get:

$$\frac{k}{y} = \frac{\alpha^2(1+g)}{\frac{1+g}{\beta} - 1 + \delta} \tag{29}$$

The first-order condition for innovation, equation **(17)**, implies:

$$\frac{R}{y} = \frac{\alpha\phi\mu\left((1+\gamma)^{1-\alpha} - 1\right)(1+g)^{\alpha}}{(1+i_L)} \tag{30}$$

The banks' zero-profit condition, equation **(18)**, implies:

$$(i_L - i) = \theta \tag{31}$$

Finally, the aggregate resource constraint, equation **(20)**, implies:

$$1 = \frac{c}{y} + \frac{I}{y} + (1+\theta)\frac{R}{y} \qquad \textbf{(32)}$$

4.2 Data and calibration

Table A: Calibrated parameter values

Parameter	Value	Comment
α	0.3133	Set to match the labour share in the UK data
β	0.9983	Set to match the risk-free rate in the UK data
δ	0.0118	Set to match the investment to capital ratio in the UK data
θ	0.0043	Set to match the spread of loan to risk-free rates in the UK data
g	0.0050	Set to match the growth rate in the UK data
μ	0.5	Arbitrary
ϕ	0.99	Set as high as seemed reasonable
σ_h	2.3256	Inverse Frisch elasticity of labour supply (Harrison and Oomen 2010)

Our calibrated parameter values are shown in Table A. We can think of there being two types of parameters: those that govern the steady state and those that govern the model's dynamics.

Starting with those parameters that govern the steady state, we assume that one period in the model represents one quarter. We first set the growth rate of the economy, g, to 0.005, implying an annual rate of growth equal to the average annual rate of growth of real consumption expenditure, *ABJR+HAYO*, over the period 2000-2008. We then use the steady-state version of the IS curve (Equation **22**) to set the discount rate, β. That is, we set β equal to $\frac{(1+g)(1+\pi)}{1+i} = 0.9983$, where the numerator is the gross growth rate of nominal consumption (π is the average rate of inflation for the consumption deflator) – given by the average growth rate *ABJQ+HAYE* over the period 2000-2008 – and the denominator is the gross nominal risk-free interest rate – given by the average Bank of England Policy Rate, *AMIH*, over the period 2000-2008. The depreciation rate for capital is related to the investment to capital ratio as shown in Equation **(27)**. Using business investment, *NPEK*, and the capital services series whose construction is described in Oulton and Srinivasan (2003), we found an average ratio of investment to capital over the period 2000-2008 of 0.0165. This value, together with a growth rate of 0.005, implies a value for δ of 0.012. Equation **(30)** enables us to set θ to match the average spread of loan rates for UK Private Non-Financial Corporations (PNFCs), *HSDC*, over the risk-free rate of 1.72% per annum over the period 2000-2008; that is, we set θ to 0.0043.

In the absence of any information, we set the probability of a successful innovation in steady state, μ, to 0.5. We set the elasticity of this probability with respect to R&D spending, ϕ, to 0.99. A value of unity meant that the Blanchard-Kahn conditions were not satisfied and that there was no stable solution for the model. Values greater than unity imply increasing returns to

R&D spending with respect to the probability of innovating; we found this implausible. Setting this elasticity to 0.99 implies a ratio of R&D spending to the sum of consumption and investment (equivalent to GDP in our model) equal to 0.1% as compared with 1.5% in the UK data.

Finally, we set α so as to match the labour share:

$$\frac{wh}{GDP}=\frac{w}{c+I}=\frac{w/y}{1-(1+\theta)R/y}=\frac{(1-\alpha)}{\left(1-(1+\theta)\frac{R}{y}\right)} \quad \textbf{(33)}$$

The labour share in the United Kingdom (measured as total compensation of private sector workers (including self-employed) divided by private sector value added at basic prices, both in current prices) averaged 69.7% over the period 2000-2008. This implies a value for α of 0.3133.

In terms of parameters governing the dynamics of the model, we simply need to set a value for the (inverse) Frisch elasticity of labour supply. Harrison and Oomen (2010) estimated the Frisch elasticity to equal 0.43 using UK data, suggesting that labour supply is inelastic. In terms of our parameter, this implies a value for σ_h of 1/0.43, that is, 2.3256, and this is the value we use.

Finally, in order to gauge the effects of a financial sector shock on the economy, we need to specify a process for this shock. We can recall that equation **(18)** implies the shock, θ_t, to be equal to the spread between the loan rate for R&D spending and the risk-free rate. So, we again use data on the loan rate for UK PNFCs together with the Bank of England base rate to construct a series for this spread. Our series is plotted in Chart 5, below.

Chart 5: Spread of PNFCs loan rates over Bank Rate

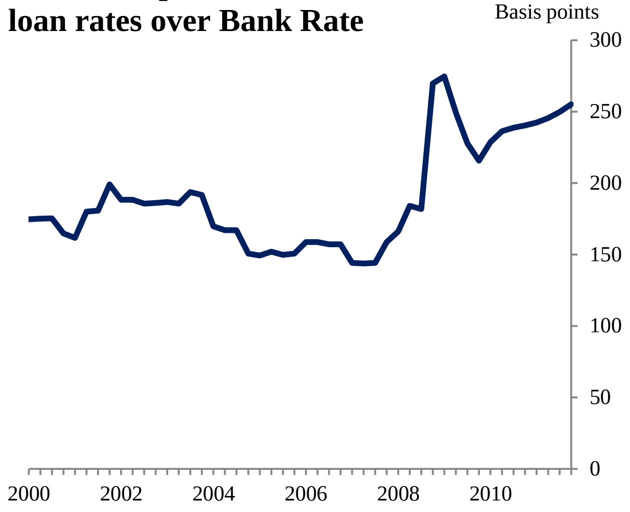

Between 2000 and 2008 Q3, the spread hovered around 150-200 basis points. But in 2008 Q4, the quarter during which Lehman Brothers went bankrupt, the spread jumped to 269 basis points and has remained high since. In 2011 Q4, it still stood at 255 basis points. Since our aim is to examine the effects of the financial crisis, we assume that agents were surprised by it. In particular, we estimate an AR(1) process for the financial shock over the period 2000 Q1 - 2008 Q3 and assume that agents expect this process to continue over the financial crisis period, that is, they do not learn that the shock turns out to be more persistent. Our estimated process for the shock is:

$$\theta_t = 0.0006 + 0.86\theta_{t-1} + \nu_t \qquad \textbf{(34)}$$

where ν is a white noise shock with a standard deviation of 0.0002 (equivalent to 8 basis points at an annualised rate).

Given our calibration of the model and the estimated process for our financial shock, we next examine the properties of our model before considering its ability to explain what we saw in response to the financial crisis.

5 Model properties

In order to assess the properties of our model, we first log-linearise the first-order conditions. As this is an endogenous growth model, our growing variables – specifically, consumption, investment, capital stock, output, R&D spending and the real wage – will be driven by a single stochastic trend. So, we first detrend these variables by dividing through by TFP, *A*, before taking the log-linear approximations of the detrended variables around their (now non-stochastic) steady-state values. We then use standard techniques to solve for the decision rules of our endogenous (detrended) variables as a function of the state variables – lagged (detrended) capital, lagged (detrended) R&D spending and lagged TFP growth – and the current shock to the spread. Given a path for TFP, and decision rules for the detrended variables, we can then obtain paths for our original variables by retrending.

The log-linearised equations of the model are as follows:

IS curve: $\hat{c}_t = E_t(\hat{c}_{t+1} + \hat{g}_{t+1}) - (i_t - i)$ **(35)**

Where $\hat{g}_t = \ln\left(\frac{A_t}{A_{t-1}}\right) - g$ is the deviation of TFP growth from its steady-state value.

Labour supply: $\sigma_h \hat{h}_t + \hat{c}_t = \hat{w}_t$ **(36)**

Production function: $\hat{y}_t = \alpha(\hat{k}_{t-1} - \hat{g}_t) + (1-\alpha)\hat{h}_t$ **(37)**

Labour demand: $\hat{y}_t - \hat{h}_t = \hat{w}_t$ **(38)**

Growth of TFP: $\hat{g}_t = \frac{g}{1+g}\hat{\mu}_t$ **(39)**

Demand for intermediates: $\hat{p}_t = \hat{y}_t + \hat{g}_t - \hat{k}_{t-1}$ **(40)**

Capital accumulation condition: $\hat{k}_t = \frac{1-\delta}{1+g}(\hat{k}_{t-1} - \hat{g}_t) + \left(1 - \frac{1-\delta}{1+g}\right)\hat{I}_t$ **(41)**

Probability of successful innovation: $\hat{\mu}_t = \phi\hat{R}_{t-1}$ **(42)**

Demand for capital: $E_t\left(\hat{y}_{t+1}+\hat{g}_{t+1}\right)-\hat{k}_t \;=\frac{1}{i+\delta}\left(i_t-i\right)$ **(43)**

Demand for R&D spending: $\left(i_{L,t}-i_l\right)=E_t\left(\hat{y}_{t+1}+\alpha\hat{g}_{t+1}+\hat{\mu}_{t+1}\right)-\hat{R}_t$ **(44)**

Zero-profit condition for banks: $\left(i_{L,t}-i_l\right)-\left(i_t-i\right)=\left(\theta_t-\theta\right)$ **(45)**

Aggregate resource constraint: $\hat{y}_t=\frac{c}{y}\hat{c}_t+\left(1-\frac{1-\delta}{1+g}\right)\frac{k}{y}\hat{I}_t+\frac{R}{y}\left(\left(\theta_t-\theta\right)+\left(1+\theta\right)\hat{R}_t\right)$ **(46)**

We solved and simulated the model using *dynare* in order to examine the model's predictions for various second moments. Table B shows the standard deviations and first-order autocorrelations of consumption growth, investment growth, output growth, GDP growth – where GDP is defined as consumption plus investment – the growth rate of R&D spending, capital utilisation, TFP growth, the probability of innovating, the risk-free interest rate, the loan rate and the spread, together with their correlations with GDP growth. We can note that since the model is one of endogenous growth, the standard deviations of the levels of consumption, investment, GDP, R&D spending, output and TFP are undefined.

The results in Table B suggest that volatility in the spread between the loan rate and the risk-free rate (the only source of volatility in this model) leads to a large amount of volatility in R&D and, as a consequence, the degree of innovation in the economy. Perhaps surprisingly, this translates into little volatility in TFP growth and, hence, output and consumption, growth.

Table B: Stylised facts implied by the model

	Standard deviation (per cent or percentage points)	Persistence (AR(1) coefficient)	Cyclicality (Correlation with GDP growth)
Consumption growth	0.0109	0.0211	-0.2338
Investment growth	0.1755	-0.0894	0.7874
GDP growth	0.0100	0.4898	1
R&D spending growth	1.9176	-0.1219	-0.6332
Output growth	0.0100	0.4898	1
TFP growth	0.0122	0.7020	0.9489
Probability of innovating	2.4653	0.7020	0.9489
Risk-free interest rate	0.0011	0.9856	0.5250
Loan rate	0.0383	0.8545	-0.5998
Spread	0.0392	0.8600	-0.6013

In terms of persistence, persistence in the driving process (ie, the spread) translates into persistence in interest rates, the probability of innovating and TFP growth. Interestingly, the growth rates of consumption and investment inherit no persistence, though GDP and output growth are more persistent than their components. Finally, we can note that rises in the spread, as expected, are associated with falls in TFP growth and the probability of successfully innovating. This means that, in turn, rises in the spread will be associated with falls in the GDP output and investment growth rates.

More insight into these correlations can be obtained by examining the responses of the endogenous variables to a shock to the spread. Chart 6 shows the response of R&D spending growth, investment growth and GDP growth to a 100 basis point rise in the spread between the annualised lending and risk-free rates. Chart 7 shows the response of the probability of a successful innovation to a 100 basis point rise in the spread and the response of the spread itself.

Chart 6: Responses to a 100 basis point rise in the spread

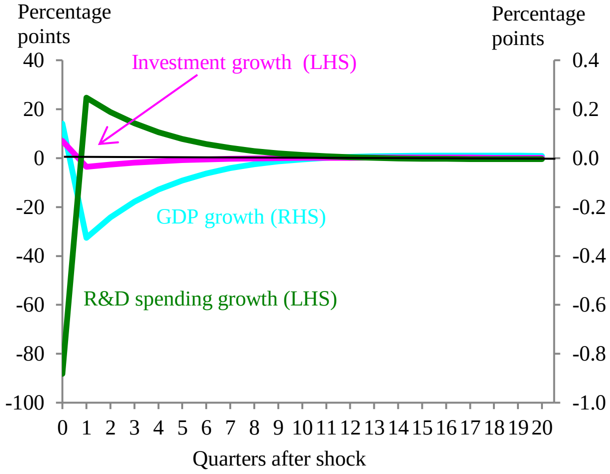

Chart 7: Responses to a 100 basis point rise in the spread

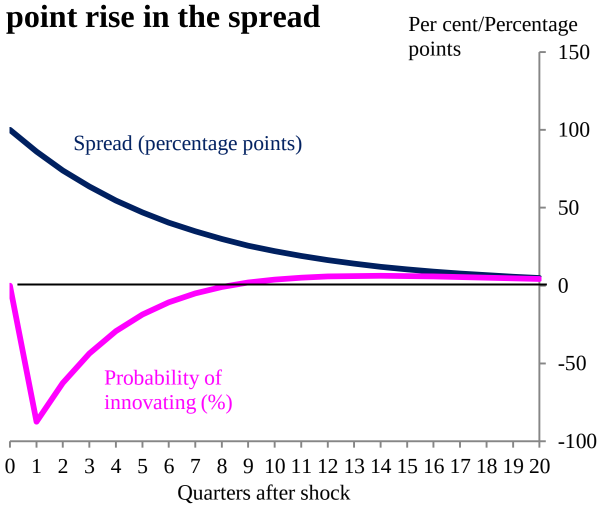

A shock to the spread leads firms to shift out of R&D spending into investment spending, given that it is relatively cheaper. This increase in investment raises GDP initially in response to the shock (though only by 0.14%). But the lack of R&D spending in the initial period ensures that innovation falls and, as a result, TFP falls in the quarter after the shock and onwards. In turn, this means that output (and GDP) falls from that point onwards. At the same time, with the spread now falling, R&D spending rises.

But the key question for this paper is to what extent the financial crisis operating through this channel can explain the fall in productivity that we have seen since 2008. Chart 8 suggests that rise in the spread should lead to a permanent fall in the level of productivity (measured here as GDP per hour). However, this fall is small for even a 100 basis point rise in the spread: less than 1%. The question is whether the financial crisis represents a large enough shock. We consider this question in the next section of the paper.

Chart 8: Response of labour productivity to a 100 basis point rise in the spread

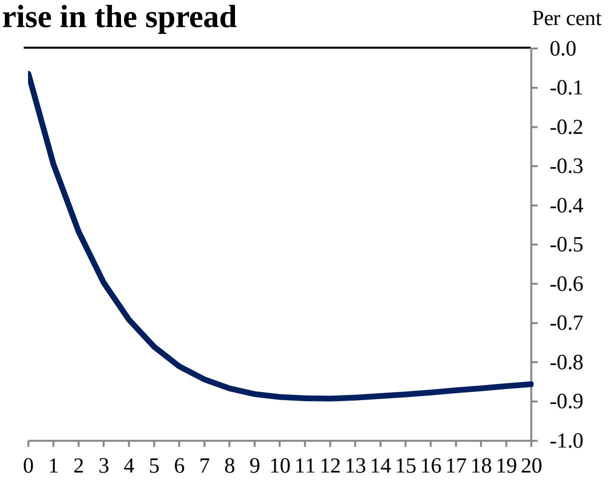

6 Can the model explain the effects of the financial crisis on productivity?

A motivation for this paper was the large fall in productivity seen in the United Kingdom in the wake of the financial crisis. Given our model, it would seem worth asking to what extent it suggests that the financial crisis can account for the weakness in UK productivity we have seen 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 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. That said, we still feel that the exercise is worth doing as it can give us a steer as to the extent to which the financial crisis might have reduced innovation and permanently damaged the level of output and productivity, an important question to be able to answer in the monetary policy context.

We perform the following simple experiment using our model. We first construct a series for the 'financial shock', ν, that replicates what happened to spreads 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 outturns with the UK data on labour productivity.

Chart 9 shows the financial shock implied by the UK data on the spread shown in Chart 1 and the estimated shock process, equation **(37)**. As can be seen the Lehman Brothers collapse in 2008 Q4 is picked up as an extremely large shock to the spread: equal to 11 standard deviations. Although the spread then fell as expected for a while, in 2010 through 2012 it has stayed higher than expected; this is picked up by the model as a sequence of positive financial shocks.

Chart 9: Financial shock

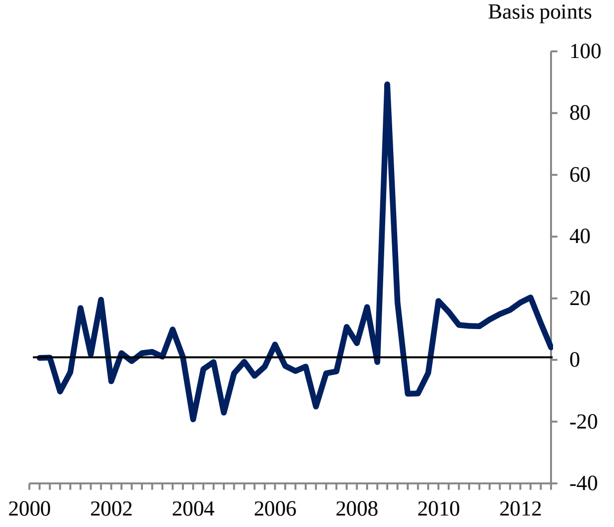

We now carry out the following experiment. We suppose the economy was on trend in 2008 Q3. We then shock the economy with the sequence of financial shocks seen since then (as shown in Chart 9). Chart 10 plots the response of output, total hours, labour productivity and TFP predicted by the model. Chart 11 plots the response of labour productivity growth and compares this with labour productivity growth in the United Kingdom – measured by the growth rate of real GDP divided by employment demeaned by its 2000-2008 average – over the same period.

The model suggests that we might expect the financial shock to lead to falls in GDP, TFP and labour productivity and Chart 11 suggests 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.

Chart 10: Simulated data

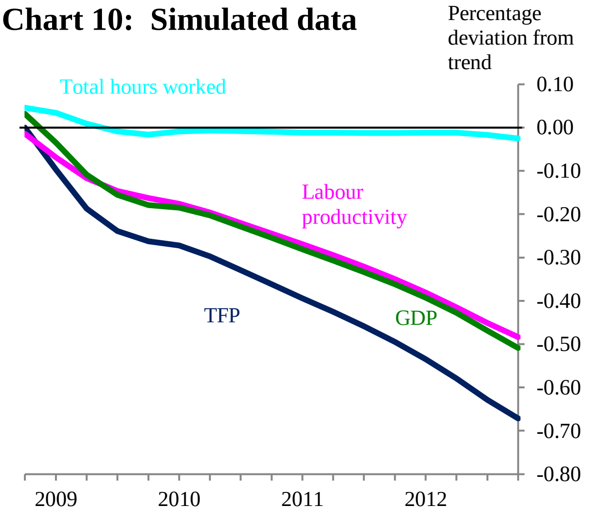

Chart 11: Labour productivity growth

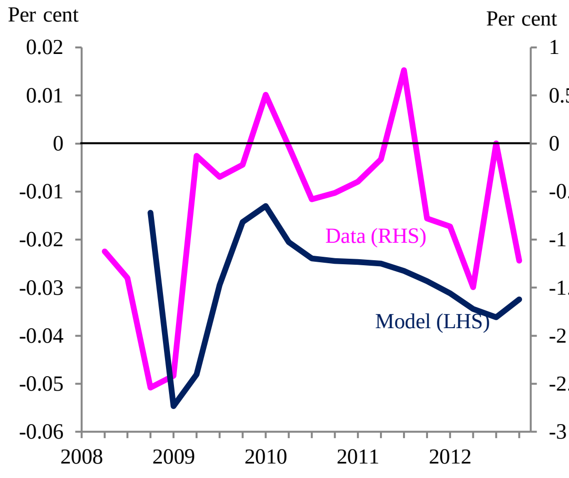

The obvious question is why is the productivity response to the financial shock operating through this channel so small?

First, there is the way we model the shock itself. Although we take as our starting point the idea that the spread can act as a sufficient statistic for the efficiency of the financial sector, it is more than possible that it simply does not capture everything. In particular, our approach suggests that the financial shock at the time of the Lehman Brothers collapse was worth only about 80 basis points on the spread and was temporary (as agents expected spreads to come back down as slowly as they did to begin with). But, the shock was also associated with a collapse in bank lending: that is, firms were unable to obtain credit at any price. To the extent that these quantity constraints were more binding (once allowing for the rise in spreads) than they had been in the past, then our approach would have underestimated the effect of the financial crisis.

Then there is the importance of R&D spending within the economy. The steady-state share of R&D in GDP implied by our model is 0.1% as opposed to the 1.5% it is in the UK data. That is, it is about 15 times larger in the data than in our model. As a result, it is quite likely that our model is going to significantly understate the importance of any shock to R&D spending on productivity. In order to resolve this problem, we need to find a way of increasing the benefits of R&D spending for firms so as to encourage them to spend more on R&D in steady state. One way of doing so might be to increase the length of time over which firms gain from innovating, say by increasing their TFP growth rate relative to non-innovating firms for more than one period or by increasing their probability of innovating further in future periods. We aim to consider these modelling choices in future work.

The next question is therefore whether the model can account for the fall in innovation that we saw in the United Kingdom in response to the shock. Chart 4 suggests that innovation fell by 20% between 2009 and 2011. Although the survey responses possibly exaggerate the size of this fall it is still much larger than we would have expected given the model and the 6% fall in R&D spending seen in the UK data. Specifically, given our (imposed) elasticity of innovation with respect to R&D spending of 0.99, we would expect a 6% fall in R&D spending to be associated with a roughly 6% fall in innovation. This suggests that at least some of the explanation for the quantitatively small response of productivity to the financial shock in our model results from our model being unable to explain the fall in innovation seen in the UK data, possibly because the general increase in uncertainty about demand acting as a disincentive to innovation since the crisis is not captured in the model. If we put through a fall in innovation of 20% in our model, we wind up with labour productivity growth roughly 0.1 percentage points per quarter below trend. This compares with roughly 0.5 percentage points per quarter in the data. In other words, given the fall in innovation, our model would explain roughly 20% of the lower-than-expected UK labour productivity growth since the financial crisis.

What would explain the remaining 80%? There are many other aspects of the financial crisis, and the response of the economy to it, missing from this model. First, the model does not consider the demand aspects of the financial shock. It assumes that consumers can borrow at the risk-free rate; a major effect of the financial crisis was that risk premia rose on loans to consumers and quantity constraints also meant that it became harder for consumers to borrow at all. In addition, it assumes that investment is financed out of internal funds and that firms discount the net returns from investment at the risk-free rate. Again, the financial crisis led to a rise in risk premia and credit rationing to firms that had a large effect on investment. The falls in consumption and investment would, in turn, lead to a larger fall in output than implied by the model, which ignores these effects. The fall in investment would also lead to a further fall in labour productivity (though not TFP) via its effect on the capital stock. We intend to allow for demand shocks in future work, though we would not expect such shocks to imply a permanent fall in productivity other than through the investment lowering capital stock channel.

Second, although the model does contain a 'working capital' channel, the effect of the financial crisis on output working through this channel will not be picked up by the model since it is assumed that working capital borrowing is financed at the risk-free rate. Fernandez-Corugedo *et al.* (2011) show that if working capital borrowing is financed at the (risky) loan rate, then the rise in spreads seen in response to the financial crisis would have a significant, though temporary, effect on productivity via this channel. The problem here is that, in this model, competition among banks would drive the working capital loan rate down to the risk-free rate, since working capital loans are not risky. We intend to revisit this issue in future work.

And finally there are other explanations for the poor productivity performance of the United Kingdom over the past few years. We would have expected to see a temporary effect of labour hoarding over the recession and of a labour supply response to the recession. The model also left out the direct contribution of the financial sector to UK productivity, which obviously fell as a result of the financial crisis, 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. It may well be that the continued falls in labour productivity we have seen in 2011 and 2012 are picking up the delayed response of the economy to the fall in innovation we saw between 2009 and 2011, whereas the fall in productivity in 2009 was much more demand related (and so temporary).

7 Conclusions

In this paper, we have constructed a simple model with which we can identify one mechanism through which the financial crisis might have impacted TFP and used it to assess the implications for the economy. 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 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. However, the response of productivity implied by the model was very small, suggesting that this channel cannot explain much of the falls in labour productivity growth, R&D spending and innovation that we have seen in the United Kingdom since 2008. In future work we will incorporate the effects of the financial crisis on demand, model better the 'working capital' channel, and look harder at the links between the financial crisis, R&D spending and innovation within the model.

References

Aghion, P and Howitt, P (2000), *The Economics of Growth*. Cambridge, MA: MIT Press

Barlevy, G (2007), 'On the cyclicality of research and development', *American Economic Review*, Vol. 97, pages 1,131-64.

Barnett, A, Broadbent, B, Chiu, A, Franklin, J and Miller, H (2014a), 'Impaired capital reallocation and productivity', *National Institute Economic Review*, No. 228, pages R35-R48.

Barnett, A, Chiu, A, Franklin, J and Sebastiá-Barriel, M (2014b), 'The productivity puzzle: A firm-level investigation into employment behaviour and resource allocation over the crisis', Bank of England *Working Paper* No. 495.

Bernanke, B and Gertler, M (1989), 'Agency costs, net worth and business fluctuations', *American Economic Review*, Vol. 79, pages 14-31.

Bernanke, B, Gertler, M and Gilchrist, S (1999), 'The financial accelerator in a quantitative business cycle framework', in Taylor, J B and Woodford, M (eds.) *Handbook of Macroeconomics*. Amsterdam: Elsevier.

Caballero, R J, Hoshi, T and Kashyap, A K (2006), 'Zombie lending and depressed restructuring in Japan', National Bureau of Economic Research *Working Paper* No. 12,129.

Comin, D and Gertler, M (2003), 'Medium-term business cycles', Starr Center for Applied Economics, New York University, *Working Paper* No. 03-05.

Corrado, C, Hulten, C and Sichel, D (2005), 'Measuring capital and technology: An expanded framework', in National Bureau of Economic Research *Measuring Capital in the New Economy*, Chicago: University of Chicago Press.

Corrado, C, Hulten, C and Sichel, D (2009), 'Intangible capital and US economic growth', *Review of Income and Wealth*, Vol. 55, No. 3, pages 661-85.

Dal Borgo, M, Goodridge, P, Haskel, J and Pesole, A (2013), 'Productivity and growth in UK industries: An intangible investment approach', *Oxford Bulletin of Economics and Statistics*, Vol. 75, pages 806-34.

Dekle, R and Kletzer, K (2003), 'The Japanese banking crisis and economic growth: Theoretical and empirical implications of deposit guarantees and weak financial regulation', *Journal of the Japanese and International Economies*, Vol. 17, pages 305-35.

Dekle, R and Kletzer, K (2004), 'Deposit insurance, regulatory forbearance and economic growth: Implications for the Japanese banking crisis', CES Ifo *Working Paper* No. 1,136.

Disney, R, Haskel, J and Heden, Y (2003), 'Restructuring and productivity growth in UK manufacturing', *Economic Journal*, Vol. 113, pages 666-94.

Fernandez-Corugedo E, McMahon, M, Millard, S and Rachel, L (2011), 'Understanding the macroeconomic effects of working capital in the United Kingdom', Bank of England *Working Paper* No.422.

Harrison, R and Oomen, Ö (2010), 'Evaluating and estimating a DSGE model for the United Kingdom', Bank of England *Working Paper* No. 380.

Jaimovich, N and Floetotto, M (2008), 'Firm dynamics, markup variations and the business cycle', *Journal of Monetary Economics*, Vol. 55, pages 1,238-52.

Kiyota, K, Nakajiama, T and Nishimura, K (2005), 'Does the natural selection mechanism still work in severe recessions? An examination of the Japanese Economy in the 1990s', *Journal of Economic Behavior and Organization*, Vol. 58, pages 53-78.

Kiyotaki, N and Moore, J (1997), 'Credit cycles', *Journal of Political Economy*, Vol. 99, pages 220–64.

Lee, Y and Mukoyama, T (2008), 'Entry, exit, and plant-level dynamics over the business cycle', Federal Reserve Bank of Cleveland *Working Paper* No. 718.

Michalopoulos, S, Laeven, L and Levine, R (2009), 'Financial innovation and endogenous growth', National Bureau of Economic Research *Working Paper* No. 15,356.

Nuño, G (2011), 'Optimal research and development and the cost of business cycles', *Journal of Economic Growth*, Vol. 16, pages 257-83.

Oulton, N and Sebastiá-Barriel, M (2013), 'Long and short-term effects of the financial crisis on labour productivity, capital and output', Bank of England *Working Paper* No. 470.

Oulton, N and Srinivasan, S (2003), 'Capital stocks, capital services and depreciation: An integrated framework', Bank of England *Working Paper* No. 192.

Reinhart, C M and Rogoff, K S (2009), *This time is different: Eight centuries of financial folly*, Princeton: Princeton University Press.

Schumpeter, J A (1942), *The theory of economic development*, Cambridge, MA: Harvard University Press.