How much does bank capital matter?

By David Aikman and Gertjan Vlieghe of the Bank's Monetary Assessment and Strategy Division.

In this article we consider how the composition of banks' balance sheets between capital and deposits affects the transmission of economic shocks. We use a small, stylised model of the economy to analyse under which conditions firms are unable to borrow as much as they would like from banks, and banks are unable to attract as many deposits as they would like from households. We show that, following shocks to aggregate productivity and bank net worth, the response of output in this model economy with credit constraints is both larger and longer-lasting than in a similar economy where credit constraints do not bind. This is because an adverse shock lowers bank capital, which constrains lending to firms and amplifies the fall in output; and it takes time for banks to rebuild their capital so it takes time for output to return to its initial level. We find that, in our model, only a small proportion of the fluctuations of output in response to productivity shocks is due to the bank capital channel, but this channel is more important when there are direct shocks to bank capital.

Introduction

The role of banks in the economy continues to be of interest to both policymakers and academics. The most recent illustration of this is the economic situation in Japan. The health of the banking sector is a key feature of most analyses of the Japanese economy.(1) Elsewhere too, concerns have occasionally been raised over banking sector health and its effect on the wider macroeconomy.(2) In this article, we investigate the role of bank balance sheets in the economy.(3) In particular, we ask whether weak bank balance sheets are likely to make an economy more vulnerable by causing a greater contraction in loans and, ultimately, output, in response to adverse shocks compared with an economy with healthy banks or an economy where banks play a less important role in financing investment. Our focus is on the relative impact of different shocks, and on the underlying factors in the economy that determine whether the impact of a particular shock is exacerbated because of bank balance sheet positions.

Why do banks and their balance sheets matter?

Banks fulfil a variety of roles in the economy: among other things, they provide payment services, transform liquid short-term deposits into illiquid long-term loans, and monitor borrowers on behalf of depositors. In practice, these functions are all important, and are at least partly interrelated. For example, banks are well placed to monitor borrowers precisely because they enter into long-term relationships with them, which may give them better information than the average depositor could obtain on his own about the likelihood that a borrower will be able to repay the loan. It is generally accepted that bank finance is likely to be more important for some types of borrowers than for others. Smaller firms, or firms without a sufficient credit history, are more likely to be dependent on banks to fund their investment expenditure. The importance of bank finance also partly depends on the structure and history of the domestic financial system.(4)

(1) See, for instance, Kashyap (2002), IMF (2003), and the article in the Winter 2003 *Quarterly Bulletin* by Farrant, Markovic and Sterne (2003).
(2) In recent years, this issue has arisen in both developed and developing countries. During the late 1980s and early 1990s, for instance, several OECD countries (Scandinavia and the United States in particular) experienced banking crises that coincided with contractions in output. At the time, many observers attached a causal role to developments in the banking sector. Similarly, during the Asian crisis of 1997, banking sector problems were seen by a number of observers as playing a causal role in the downturn. Hoggarth *et al* (2002) provide a cross-country comparison of the output cost of several recent banking crises.
(3) We build on a recently developed theoretical framework by Chen (2001). The particular version of the model we use differs in a few details, as described in the appendix.
(4) The evolving role of banks is discussed in a speech by former Bank of England governor Sir Edward George (1997).

Given that banks are important, why does the structure of their balance sheets matter? Table A shows a simplified version of a bank balance sheet. Banks hold a variety of assets: they make loans to consumers, to non-financial companies, to other financial institutions, and to governments. Such loans can be made with an explicit loan agreement (eg a mortgage), or by buying a bond issued by a company or the government. These loans are financed by deposits and bank equity, which together represent banks' liabilities. Deposits can be in the form of accounts held at the bank, or in the form of debt issued by the bank, such as certificates of deposit, or longer-maturity debt instruments. The equity portion of the liabilities consists of money injected by shareholders, and profits retained by the bank. For the purposes of this paper, we will call the equity portion of liabilities 'bank capital'.(1)

Table A
A simplified bank balance sheet

Assets	Liabilities
Loans	Deposits Equity

A celebrated insight by Modigliani and Miller (1958) is that, if firms have perfect access to borrowing markets and there are no other distortions in the economy (such as taxes, bankruptcy costs or imperfect information), the structure of liabilities will not affect their investment decisions. Interpreting banks as 'firms', this suggests that a bank's lending decision will be unaffected by the structure of its balance sheet. Why might this logic fail to hold in practice, so that the quantity of bank capital does influence lending decisions?

Two elements are required: first, it must be more costly for banks to raise equity finance rather than deposit finance. That would, of course, lead banks to issue only the lowest cost finance, ie deposits. So a second element is required: there must be some reason why banks need to hold capital. Equity finance may be more costly than deposit finance because of transaction costs(2) or taxes (if interest payments are tax deductible and dividend payments are not). Some equity finance may nevertheless be required for the following reasons. First, bankruptcy is not costless, so a sufficient 'buffer' of equity may be needed to protect a bank against unanticipated losses and reduce the cost of debt. Second (and related to the previous point), capital regulation generally requires banks to hold some minimum level of equity as a fraction of their assets.(3) Third, informational problems may exist between the banks and their depositors, which can be alleviated if the bank holds sufficient amounts of equity. For these reasons, shocks that lower bank equity may feed through into lower loan supply.

This article focuses on informational problems as the reason why banks are required to inject equity. In particular, we assume that depositors cannot observe how much risk a bank is taking in its lending business. But depositors know that banks with more equity invested in the lending will be less likely to take excessive risks—equity, after all, represents the shareholders' own money that is at stake, so they have more to lose if things go wrong. So depositors may only be willing to put their money in a bank with sufficient equity. The amount of equity therefore limits how much deposit finance a bank can attract, which in turn limits the amount a bank can lend.

If banks fulfil an essential role in the economy, and the structure of their liabilities matters, then weakened banks (ie those with lower capital) must lead to worse outcomes. But the more difficult question to answer is: how much worse will these outcomes be, and what do they depend on? There is much empirical evidence on this matter,(4) but the evidence is often open to alternative interpretations. This is because it is difficult to disentangle relative movements in the supply of bank loans versus the demand for loans: bank loans, bank capital and interest rates on loans can all be expected to move with the economic cycle, even if they do not themselves affect the economic cycle. They may just reflect demand conditions. A further difficulty is that, even if bank balance sheets have only small effects on average, there may be episodes when they become much more important. Such episodic effects are difficult to estimate empirically. Given these difficulties, a complementary approach is to try and give a theoretical answer to the problem: under certain assumptions about how consumers, firms and banks behave in an

(1) In practice, the definition of bank capital for regulatory purposes also includes some forms of longer-term debt issued by the bank.
(2) Interpreted broadly, these could include both direct transaction costs incurred in issuing equity, and the implicit costs resulting from the fact that equity issues are sometimes interpreted as signalling bad news about the bank's profitability.
(3) The details of these requirements are described in BIS (1988).
(4) See, for example, Peak and Rosengren (1997), Kashyap and Stein (2000), Angeloni *et al* (2003) and references therein.

economic model, how much does bank capital matter and what does it depend on? This is the approach taken in this article.

A stylised model of banks

To investigate the role of banks in the economy, we use a theoretical model by Chen (2001).[1] In this simple set-up, households deposit their savings in banks, and banks make loans to entrepreneurs (ie firms) to finance investment in physical capital.

The structure of balance sheets plays an important role in this model because it is assumed that there is imperfect information in the economy. Specifically, households cannot perfectly observe banks' lending activity, and banks cannot perfectly observe entrepreneurs' investment activity. This creates a so-called 'moral hazard' problem for both banks and entrepreneurs. The intuition is as follows: entrepreneurs can choose the riskiness of their projects. In choosing the project, the entrepreneur weighs up what he earns if the project goes well against what he loses if the project fails. If the project is successful, the entrepreneur gets the output from the project plus the value of his physical capital, less what he owes the bank in interest. If the project fails, there is no output, and no interest to pay, but he loses any of his own wealth that was invested in the project. Provided that entrepreneurs have a large enough stake (ie entrepreneurial net worth) in the project, they will not be tempted to choose excessively[2] risky projects because they will have too much to lose if the project fails.

A similar moral hazard problem occurs in the banking sector: banks can monitor entrepreneurs and thereby influence the range of projects available so that less risky projects are chosen.[3] But monitoring is costly for the banks, and depositors cannot observe whether the bank is monitoring or not. Banks may therefore be tempted not to monitor if the costs are too high. But if banks have a large enough stake (ie bank capital) in the projects they are lending to, they will always monitor, because they have too much to lose by letting entrepreneurs choose risky investment projects.

This is therefore an economy with credit constraints: the quantity of bank capital affects how much banks are able to lend and the quantity of entrepreneurial net worth affects how much entrepreneurs are able to borrow. We now consider how the economy responds over time to shocks. We first consider a shock to productivity, and then a direct shock to bank net worth.

The economy's response to a productivity shock

Chart 1 shows how the key variables in the model respond to a persistent negative shock to productivity. An adverse productivity shock initially reduces output, which lowers the *ex-post* return on all entrepreneurial projects. Bank and entrepreneurial net worth thus immediately fall.[4] In an economy where credit constraints do not bind, the effect on output would stop there.[5] But with binding credit constraints, there are second-round effects.

There are two distinct channels through which these second-round effects occur. First, the bank capital channel. With less of their own money at stake, and engaging in lending to borrowers who have lower productivity returns, banks are perceived by households as having a smaller financial stake in the projects they are monitoring, and this makes them riskier places to deposit funds. Banks therefore find it harder to attract deposits. Less capital and deposits mean fewer resources are available for lending, and loan supply contracts. Second, the borrower net worth channel. It acts in a similar way: with less net worth and lower expected returns (due to lower productivity), borrowers are viewed as having less at stake in the outcome of their projects—they have less to lose when projects fail. Banks therefore curtail their lending even further.

As a result of this squeeze on credit, entrepreneurs are able to buy less capital for use in the following period. This shift lowers expected future returns from capital, depressing the current price of capital. Moreover, part of entrepreneurs' net worth consists of their holdings of this physical capital, so the drop in its price further reduces current entrepreneurial net worth. There is

(1) Chen's model in turn builds on the work of Holmström and Tirole (1997).
(2) All possible projects that the entrepreneur can choose are risky, but it is assumed that the riskier projects have negative net present value to society, so it is not desirable for entrepreneurs to pursue them.
(3) It is assumed that banks can only partially influence entrepreneurs' choices. Since they cannot fully control which projects the entrepreneurs choose, both banks and entrepreneurs need to have enough of their own wealth at stake for the best investment projects to be chosen.
(4) In this dynamic model, banks' net worth consists of retained profits from their lending activity. And entrepreneurs' net worth consists of retained profits from their production activity plus the value of physical capital retained from the previous period.
(5) In this model, labour supply is assumed to be fixed so there is no endogenous labour supply response.

Chart 1
Response to a productivity shock

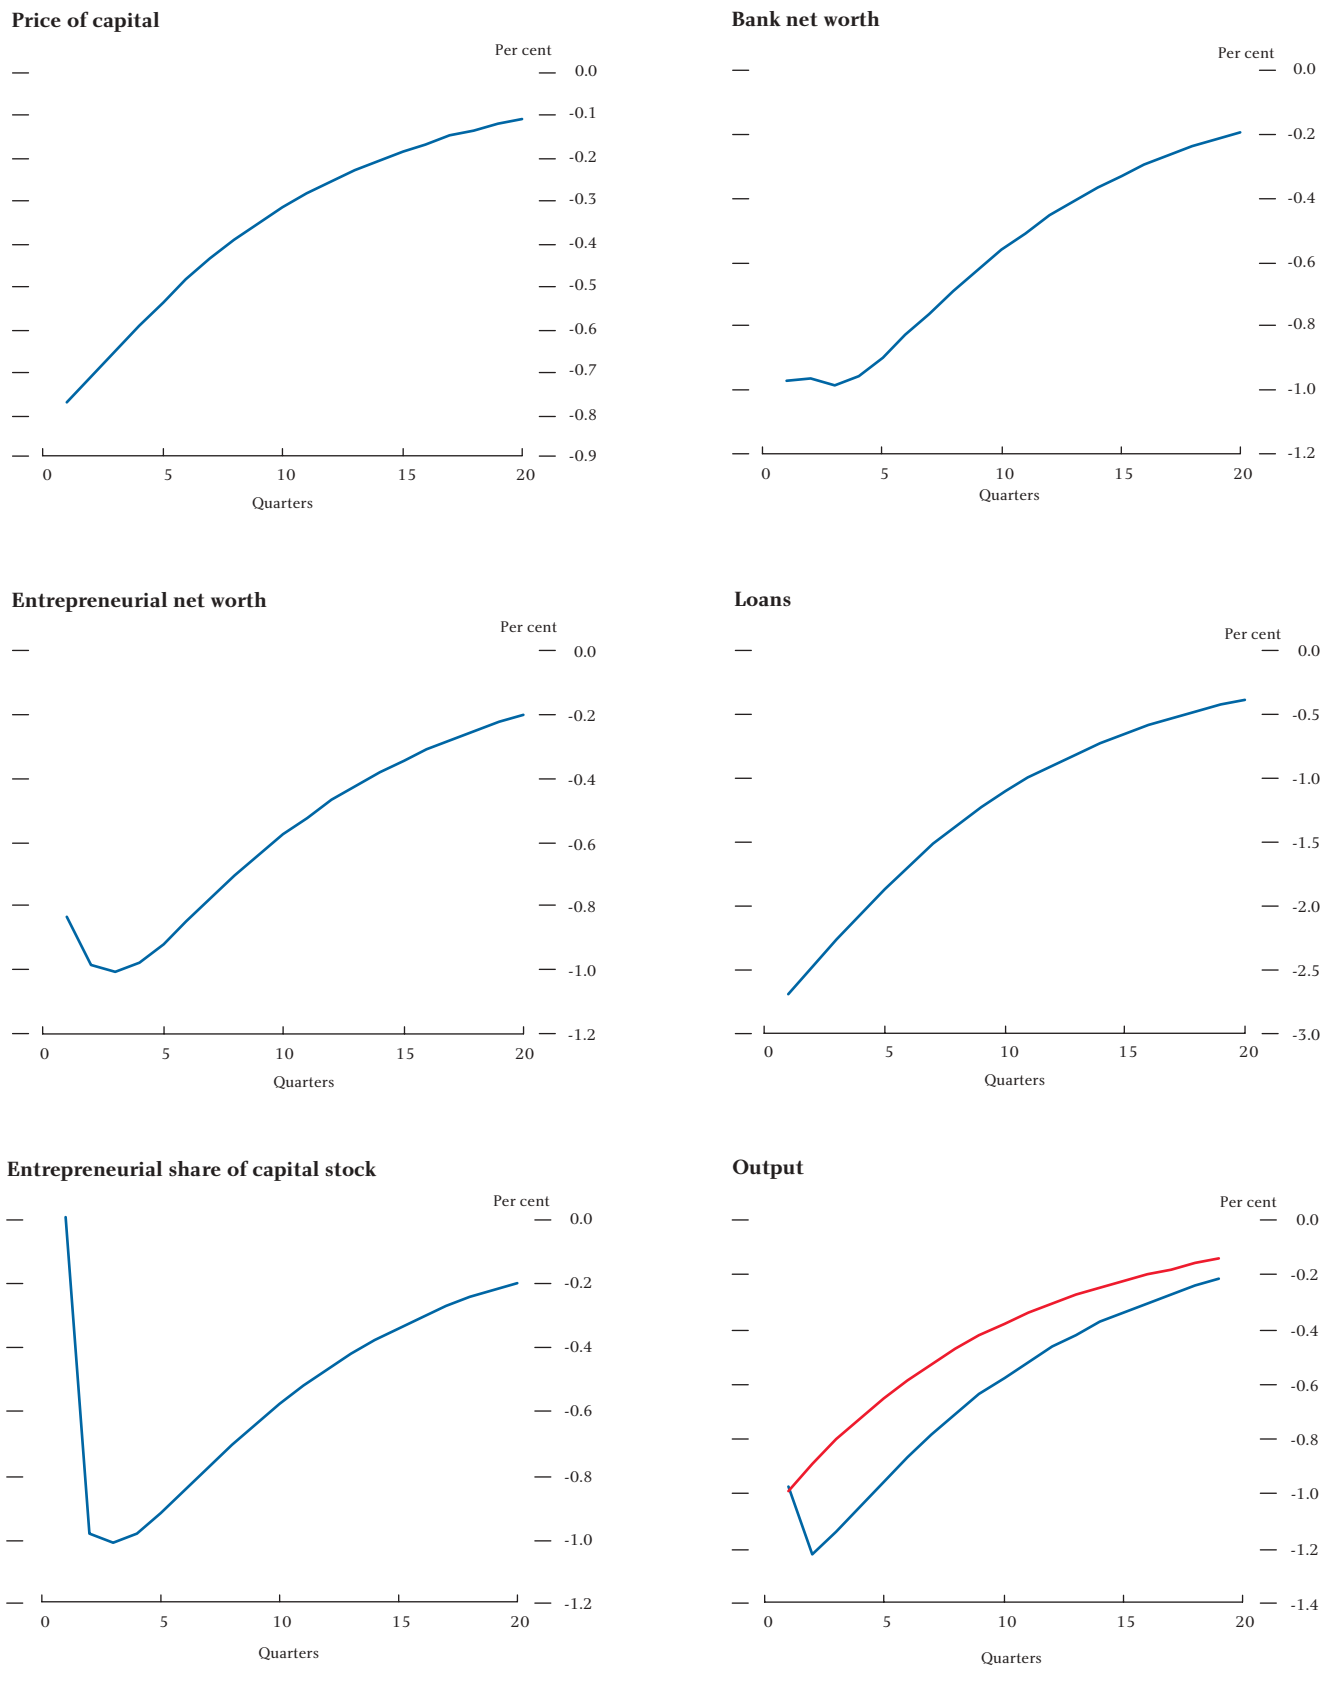

Note: Responses to a 1% fall in the level of productivity, with an autocorrelation of 0.9 (ie 90% of the shock persists into the next period and so on). Units along the vertical axis are percentage deviations from the initial level of each variable. The blue line in the bottom right panel represents the response of aggregate output when credit constraints are binding; the red line represents the output response when credit constraints are non-binding. The time scale along the horizontal axis represents quarters. And the shock occurs after one quarter.

therefore a feedback effect from net worth to capital good prices, and then back from capital good prices to net worth. And this magnifies the impact of the initial shock.

How much amplification and persistence is generated by the binding credit constraint? The bottom right panel in Chart 1 illustrates the model's prediction. It compares the output response of the constrained economy (blue line) with that of a version of the model where the borrowing constraints do not bind (red line), ie an economy where entrepreneurs can borrow as much as they need to finance profitable projects.(1) The behaviour of output is clearly quite different in the two cases. In the unconstrained economy, output starts to return to its initial level immediately following the shock. But in the constrained economy, output continues to contract for a while as the credit squeeze reduces the amount of capital that entrepreneurs can afford to buy. It takes time to close the gap between the constrained and unconstrained level of output because it takes time for banks and entrepreneurs to restore their balance sheets via the accumulation of net worth.(2)

The economy's response to a direct shock to banks' net worth

In the previous experiment, bank balance sheets acted to propagate the effect of an exogenous shock to productivity. We now analyse the effect of a direct adverse shock to bank capital; one that is entirely separate from demand and supply factors in the economy. Such a shock could be loosely interpreted as a reduction in bank capital due to a fall in the value of a bank's foreign assets, since there is no corresponding reduction in the value of domestic loans. So, if foreign banks are active in an economy, it could reflect an adverse shock to their assets in their home economy. It could also be interpreted as a one-off reduction in bank capital resulting from the discovery of fraud. Chart 2 shows the dynamic response of the economy following such a shock.

The immediate effect of the reduction in banks' net worth is to reduce the funds available for lending. This lowers the amount of capital that entrepreneurs can buy to produce output in the following period, which again lowers expected future returns leading to a fall in the current price of capital. Because the capital stock held by entrepreneurs is now worth less, their net worth also falls, which again leads to a further contraction in loan supply. And from this point on, the shock is propagated in a qualitatively similar way to the productivity shock case described earlier.

The bottom right panel of Chart 2 shows that the reduction in output following the fall in bank net worth is again persistent: it takes approximately ten quarters before output returns to its initial level. The effects of the shock are again also amplified: in an unconstrained economy, the reduction in bank net worth would have no effect on output, so the reduction in output is entirely due to credit constraints binding.

There are, however, two notable differences from the effects of the productivity shock. First, since output is initially produced with capital that was already in place, the bank net worth shock affects output with a lag rather than immediately.

Second, and more importantly, 'financial' shocks and productivity shocks affect banks' incentives to monitor—and therefore the size of the bank capital channel—in quite different ways.(3) In particular, bank moral hazard problems become less severe when bank net worth falls exogenously. So even though the reduction in bank net worth has adverse consequences, this is partly offset by a fall in the fraction of net worth that banks need to hold against loans (ie their capital to asset ratio), as Chart 3 illustrates. The opposite occurs following a shock to aggregate productivity: banks are required to take a larger stake in the projects they monitor (ie the required capital to asset ratio rises)—which reinforces the initial effect of the shock.

Why do moral hazard problems ease after a financial shock, but tighten following a productivity shock? Simply because different shocks have different effects on a bank's incentives to undertake the costly monitoring of firms that it lends to. In the case of an adverse productivity shock, the driving process is the reduction in current and future productivity of firms. This leads simultaneously to a shortage of bank capital and a reduction in the return that banks can expect on their

(1) Both responses are drawn as percentage deviations from the initial level. But note that the initial level of output in the unconstrained economy is higher as the capital stock is more efficiently used. So productivity shocks in the constrained economy cause more variability in output around a lower level.
(2) Our results about the extent to which financial variables feed back onto real variables such as output are not particularly sensitive to values chosen for the moral hazard parameters: the monitoring cost and private benefit.
(3) We focus this discussion on the effect that these shocks have on banks' incentives to monitor. But entrepreneurs' incentives to put in high effort are also affected by these shocks.

Chart 2
Response to a bank net worth shock

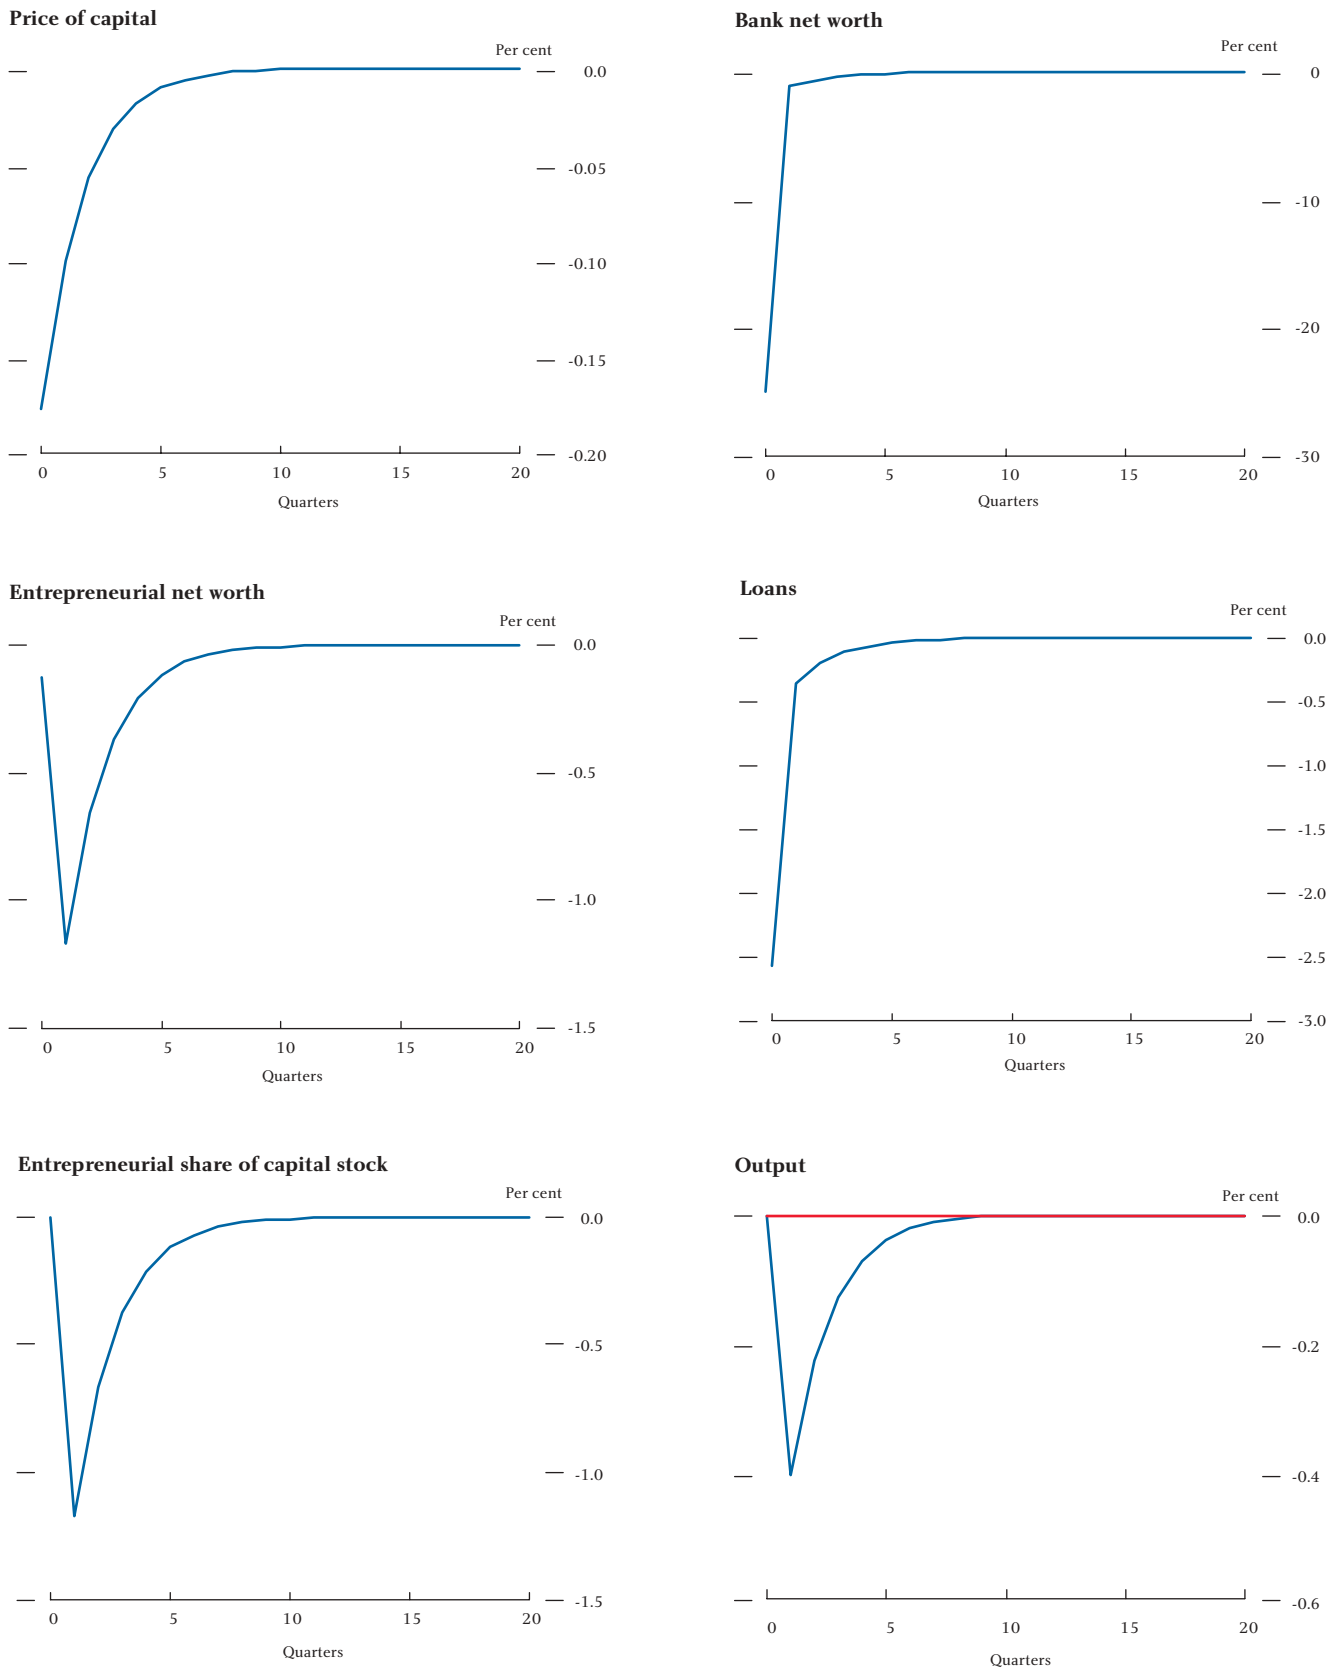

Note: The figures show the percentage deviations from long-run equilibrium for each variable following a once-and-for-all shock that reduces bank capital by 25% (ie were loans to remain constant, the capital to asset ratio of the banking sector would fall from 8%—the assumed long-run value—to 6%). The blue line in the bottom right panel represents the response of aggregate output when credit constraints are binding; the red line represents the output response when credit constraints are non-binding. The time scale along the horizontal axis represents quarters. And the shock occurs in quarter one.

Chart 3
Banks' capital to asset ratios following different shocks

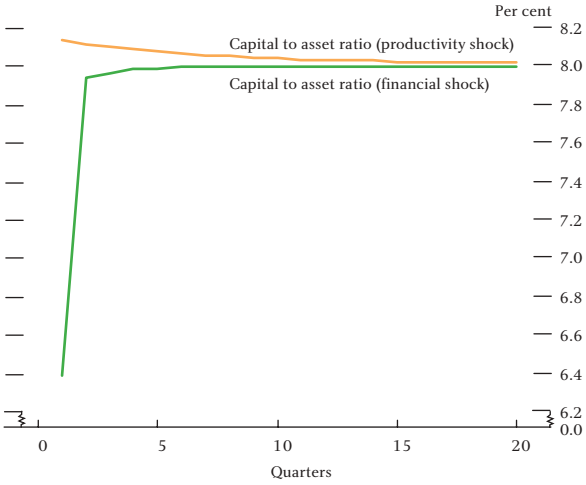

Note: Banks in our model hold capital that is worth 8% of their loans in the steady state. The units along the vertical axis represent the actual value of each ratio.

loan portfolio, because their customers, the firms, are less productive. This reduction in expected returns on loans lowers banks' incentives to behave diligently: since banks have less to lose, they have less incentive to perform the costly monitoring. Depositors realise this, and demand that banks keep a larger stake in the projects they lend to, in the form of a higher capital to asset ratio.

In the case of the adverse financial shock, the driving process is an immediate reduction in bank capital that is at first not associated with any changes in productivity. Since banks lend out their own capital as well as deposits, a reduction in bank capital reduces banks' ability to lend, so reduces the amount of financing that entrepreneurs can obtain to invest in physical capital. When entrepreneurs have lower levels of physical capital, this increases the productivity of that capital. This is the result of diminishing returns to physical capital: when firms hold more capital, the productivity of an additional unit of capital falls. So when firms are forced to hold less capital, as is the case when a financial shock reduces banks' ability to lend, the productivity of an additional unit of physical capital rises. This improves the incentives of the banks: the return they can expect from their loan portfolio has increased, so banks have a greater incentive to perform the costly monitoring. Depositors realise this, and allow the banks to operate with lower capital to asset ratios while they recover from the adverse financial shock. So while a reduction in bank capital is obviously bad for the economy as a whole, it improves the banks' incentives to act diligently, and therefore cushions the adverse effect of the shock. This contrasts with the productivity shock, where incentives worsen, and therefore amplify the effect of the shock.

These results highlight the importance of having a well specified model to analyse the role of net worth: the role of bank capital in transmitting shocks is likely to depend upon the nature of the disturbance.

How important is the bank capital channel?

What is the relative contribution of the two types of financial frictions to the amplification of shocks previously documented? To address this, we modify the model by allowing the depositors of each bank to observe (at no cost) whether banks are adequately monitoring the entrepreneurs they are lending to—effectively 'switching off' the friction that forces banks to hold capital. By comparing this economy with the benchmark economy studied above, we can then quantify the marginal role played by bank capital in this model.

We would expect a less amplified response to shocks when this friction is 'switched off': if bank capital is no longer playing a crucial role in providing the right incentives for banks, bank balance sheets cannot serve to amplify shocks. Chart 4 explores this possibility by comparing the economy's response to the adverse productivity shock considered earlier, with and without bank moral hazard. Blue lines in the figure refer to the benchmark economy studied previously; red lines refer to the economy with the bank capital channel switched off.

When the bank capital channel is switched off, there is clearly less volatility in most of the key variables of the model. As Chart 5 shows, however, the quantitative impact of this channel on the volatility of output (given by the difference between the blue and red line) is very small. The peak response of output falls only slightly(1) and there is no impact at all on its timing. Clearly, little of the amplification and persistence displayed by the benchmark economy was due to the bank capital channel.

What are the reasons for this result? The quantity of capital held by banks in the benchmark economy is very small relative to the net worth of entrepreneurs. This

(1) The peak response without bank moral hazard is 20% larger than the peak response without any constraints. It is 23% larger in the benchmark economy with both bank and entrepreneurial moral hazard, relative to the unconstrained case.

Chart 4
Responses with and without a 'bank capital channel'

Price of capital

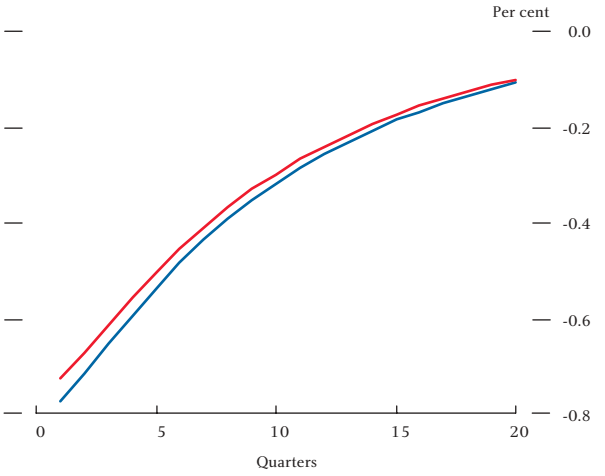

Bank net worth

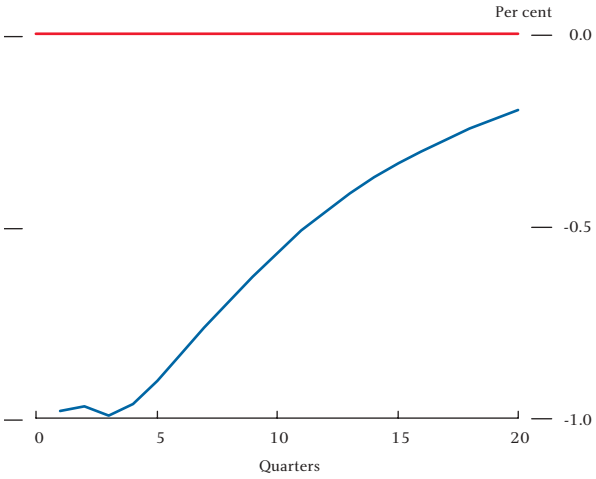

Entrepreneurial net worth

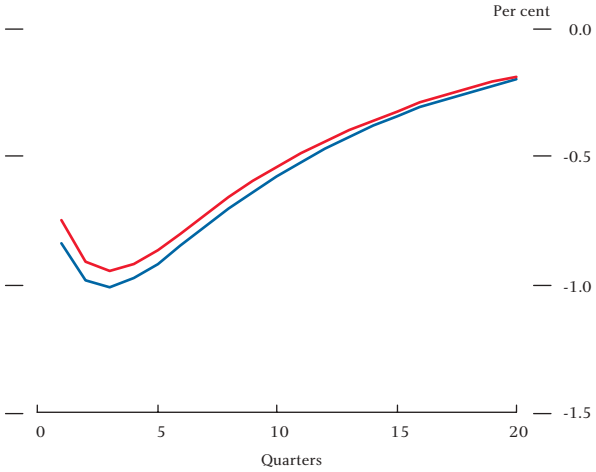

Loans

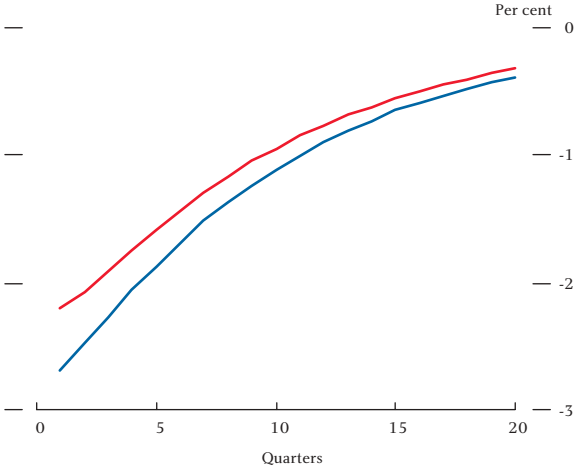

Entrepreneurial share of capital stock

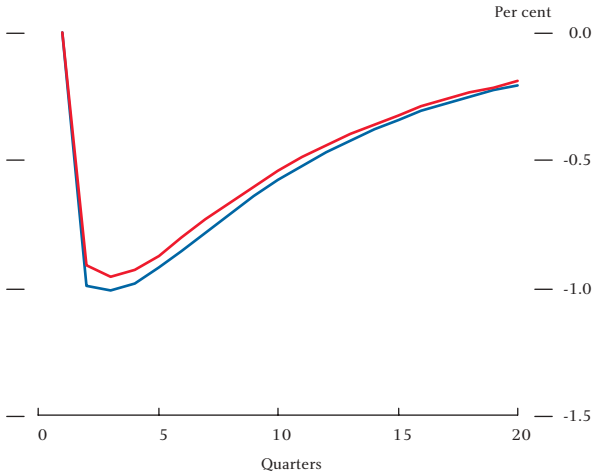

Note: Responses to a 1% fall in the level of productivity, with an autocorrelation of 0.9. Blue lines refer to the benchmark economy studied previously; red lines refer to the economy with the bank capital channel switched off. The units along the vertical axis are percentage deviations from the initial level of each variable.

Chart 5
Response of output

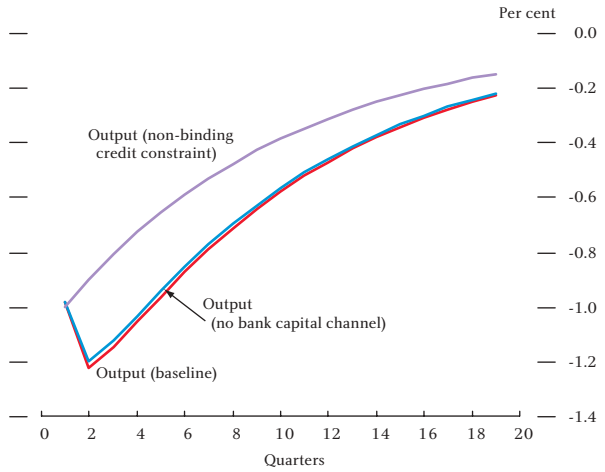

reflects the real world: firms in the United Kingdom typically hold net worth of approximately the same value as their loans; banks, on the other hand, are very highly leveraged institutions, and they typically keep a ratio of only around one tenth of the value of their loans. So entrepreneurial net worth is about ten times as large as bank net worth. If banks need to earn a larger return on their net worth as a result of increased moral hazard, this uses up only a small fraction of the total resources of the economy.

This result that bank balance sheets do not matter much in the transmission of productivity shocks should not, however, be taken to be a general conclusion. It is a consequence of the many modelling choices that have been made. It is probably related at least in part to the fact that there is no time-varying default rate for entrepreneurs in the model (firms fail in the model, but at a constant exogenous rate). So any effect of bank capital on the real economy must arise due to changes in the return that banks earn on their net worth. And as previously stressed, achieving a change in banks' returns does not require a large diversion of resources from other sectors. A legitimate concern may be that in the real world banks also lose money in recessions because a higher than expected proportion of their borrowers default on loans. Even a small change in the outstanding value of loans due to defaults can have a large impact on banks' net worth precisely because of banks' low capital to asset ratio. So the effect of a shock in such an economy might look like a combination of our productivity shock and our bank net worth shock.

Finally, we should note that eliminating bank moral hazard also affects the long-run properties of our model. Borrowing constraints ease (as the supply of loans is no longer tied to the net worth of the banking sector) and entrepreneurs end up holding a greater fraction of the capital stock. The economy therefore moves closer to the unconstrained case. Our simulations suggest that this effect is quantitatively quite large.

Conclusions

We have examined the role of bank capital in the economy using a model where asymmetric information leads to moral hazard problems between depositors and banks, and between banks and entrepreneurs. We find that the response of the economy to shocks in the presence of these financial frictions is both amplified and more persistent relative to a similar economy where credit constraints do not bind. The intuition is that an adverse shock lowers the net worth of banks and entrepreneurs, which in turn lowers the lending capacity of banks and the borrowing capacity of entrepreneurs, amplifying the effect of the initial shock on output. It takes time for entrepreneurs and banks to rebuild their net worth positions, so the effect of the shock persists. However, the nature of the shock also has important implications. Adverse productivity shocks increase moral hazard problems between depositors and banks, whereas direct shocks to bank net worth reduce these moral hazard problems.

We also find that only a small proportion of the amplification and persistence displayed by the benchmark economy is due to the bank capital channel, unless there is a direct shock to bank capital. A more general model incorporating time-varying borrower default, however, may give bank capital a more important role even for shocks that are not directly to bank capital.

Appendix

The model we use to generate the simulations in the text is based on Chen (2001). The exact version we use differs in a few details from Chen's original model as we now describe. The interested reader should consult Chen (2001) for a fuller description of the model. We use the same notation.

Our benchmark model (used to generate the impulse responses in Charts 1 to 3) contains two modifications to the basic Chen model. First, we adopt a constant elasticity of substitution specification for the marginal product of capital in home production. Denoting the technology available to households for transforming capital into consumption goods by $Y_t = G\left(K^h_{t-1}\right)$, the marginal product is defined as:

$$G'\left(K^h_{t-1}\right) \equiv \frac{1}{\beta}\left(\frac{\lambda K^e_{t-1}}{\overline{K}}\right)^{\frac{1}{\varepsilon}}$$

where β represents each agent's subjective discount factor, $K^e / \overline{K}$ represents the fraction of the total physical capital stock held by entrepreneurs, ε is the (constant) residual supply of capital from households with respect to the user cost, and λ is a scaling parameter. This affects the model properties as we move from the baseline model to the 'no bank capital channel' model.

Second, we assume that the private benefits (or reduction in effort) available to entrepreneurs when choosing riskier projects ('bad' and 'rotten' projects in Chen's terminology) are denominated in terms of consumption goods. This has the advantage that private benefits and monitoring costs are treated symmetrically in the financial contract.

To generate the results on the relative importance of the bank capital channel (Charts 4 and 5), we modify the model by making the monitoring action of banks public information (ie observable at zero cost by all other agents in the economy). This change allows us to drop banks' incentive compatibility constraint from the list of equilibrium conditions. Banks in this case will no longer be able to command a share of the surplus from projects they monitor. And like households, they will receive an expected return just sufficient to satisfy their participation constraint:

$$E_t v_{t+1} p_H R^b_t K^e_t \geq \frac{A_t}{\beta}$$

The incentive for banks to postpone consumption continually and accumulate net worth is therefore eliminated and, given linear utility, the exact time path of consumption and savings will be indeterminate. We assume that banks accumulate zero net worth and consume their endowment period by period.

Given that there is now one agent less needing to be paid off from the entrepreneurial project's surplus, *ceteris paribus* more of the surplus can be pledged to depositors. Depositors are therefore willing to invest more and leverage goes up. This effect is summarised in the following equation linking capital purchases by the entrepreneur to net worth:

$$K^e_t = \frac{W_t}{\tilde{Z}_t}, \text{ where } \tilde{Z}_t \equiv q_t + c - \beta\left[p_H\left(E_t v_{t+1} R - \frac{b}{\Delta p}\right)\right] + E_t q_{t+1}$$

References

Angeloni, I, Kashyap, A and Mojon, B (2003), 'Monetary transmission in the euro area: where do we stand?', in Angeloni, I, Kashyap, A and Mojon, B (eds), *Monetary policy transmission in the euro area*, Cambridge University.

BIS (1988), *International convergence of capital measurements and capital standards*, produced by the Basle Committee on Banking Supervision, Bank for International Settlements, Basle.

Chen, N (2001), 'Bank net worth, asset prices and economic activity', *Journal of Monetary Economics*, Vol. 48(2), pages 415–36.

Farrant, K, Markovic, B and Sterne, G (2003), 'The macroeconomic impact of revitalising the Japanese banking sector', *Bank of England Quarterly Bulletin*, Winter, pages 439–51.

George, E (1997), 'Are banks still special?', *Bank of England Quarterly Bulletin*, March, pages 113–18.

Hoggarth, G, Reis, R and Saporta, V (2002), 'Costs of banking system instability: some empirical evidence', *Journal of Banking and Finance*, Vol. 26, pages 825–55.

Holmström, B and Tirole, J (1997), 'Financial intermediation, loanable funds and the real sector', *Quarterly Journal of Economics*, Vol. 112(3), pages 663–91.

IMF (2003), *Japan Article IV Consultation—Staff report*, International Monetary Fund, Washington, D.C.

Kashyap, A (2002), 'Sorting out Japan's financial crisis', *Federal Reserve Bank of Chicago Economic Perspectives*, 4th Quarter.

Kashyap, A and Stein, J (2000), 'What do a million observations on banks say about the transmission of monetary policy', *American Economic Review*, Vol. 90(3), pages 407–28.

Modigliani, F and Miller, M (1958), 'The cost of capital, corporation finance and the theory of investment', *American Economic Review*, Vol. 48(3), pages 261–97.

Peak, J and Rosengren, E (1997), 'International transmission of financial shocks: the case of Japan', *American Economic Review*, Vol. 87, pages 495–505.