A submission to the Walker Review

Sam Robbins and Leigh Caldwell, *Intellectual Business*

30th May 2009

Intellectual Business is a think tank promoting innovative solutions to business and financial problems. Its research combines economics, psychology, theory of business and law, and an understanding of public policy, to build models of how people work and trade together – an economics based on realistic psychological models and their macro consequences.

Sam Robbins is a former Executive Director of the Equities Division of UBS Investment Bank, London. He has over ten years experience of financial markets, having worked first as a quantitative analyst and later as a senior derivatives trader. He has a BA in Mathematics from the University of Cambridge and an MSc in Government, Policy and Politics, from Birkbeck, University of London.

Leigh Caldwell is founder and director of Intellectual Business. He is also chief executive of Inon, a technology and economics consultancy founded in 1995 and focused on applying behavioural economics using Web technology. He has written for VoxEU on behavioural economics and financial regulation, and has a BSc (Hons) in Mathematics and Physics from the University of Glasgow.

Intellectual Business

23 Jacob Street
London
SE1 2BG

020 7064 6556
http://www.intellectualbusiness.org/

Introduction

In examining corporate governance of financial companies, it is reasonable to ask: why do the normal incentives of ownership not encourage appropriate behaviour? Shareholders of banks should be able to discipline excessive risk-taking and encourage high long-term returns, and competitive market discipline should ensure the optimal social return to financial activities.

In this submission we will argue that there are two main reasons why this does not occur.

The first is that the externalities of the financial sector are not priced correctly. Specifically, limited liability for shareholders means that they do not bear the full cost of the risks they take; and the short-term incentives for financial institutions do not reward the long-term return to investment, which discourages it from taking place. Furthermore, we argue that bank boards did not deviate from the wishes of their shareholders (prior to the recent crisis) but reflected exactly their revealed preferences for excessive risk.

The second is that shareholders and other financial market participants are subject to bounded rationality. There are limits on the information they can process, so they do not make fully informed decisions and can be taken advantage of by better-informed insiders; and they are subject to systematic cognitive biases which can distort behaviour and diminish long-term returns.

We explore those problems in further detail below, and then present potential solutions.

Externalities: Limited liability

In this section we argue that the ownership of UK banks by large numbers of small shareholders contributed to the culture of excessive short-term risk-taking that existed prior to 2007. A shareholder's limited liability creates an incentive for them to demand excessive amounts of leverage, which would not be demanded if they were exposed to the downside this engenders. Whereas in other industries this pressure towards excessive risk taking is counterbalanced by the existence of creditors' desire to protect their investments, we argue that large banks encountered less resistance from their creditors, because of depositor guarantees and a sector-wide interest in preserving the status quo. This led to the running down of capital ratios and the increased exposure of the sector to systemic risk. We will also argue that, though the structure of bankers' remuneration did lead to excessive long-term risk-taking for short-term gain, this culture evolved to satisfy the preferences of shareholders, and any solution to the problem of excessive risk-taking must therefore encompass how banks are owned as well as how they are managed.

The larger UK retail banks are typically owned by large numbers of small shareholders. Our experience in the financial sector suggests that investors' collective behaviour, prior to 2007, was to actively target the banks offering maximum returns on investment, without any particular attention to the risks involved. Investors would routinely target returns on equity as high as 20% and punish stocks which did not achieve such exaggerated rates of return. We argue that banks and other financial institutions increased their leverage as a response to this demand. Shareholders rationally demanded this extra risk be taken, because their limited liability meant that, whilst their downside was strictly capped at the level of their investment, their upside exposure could effectively be multiplied through the process of leveraging.

Although the limited liability of shareholders creates the potential demand for excessive risk-taking in any company, we argue that in most cases this is not realised because of the countervailing power of their creditors. Banks are an exception to this rule, as they have considerable power to determine how leveraged they are. Firstly, small depositors, who provide a significant fraction of all the liabilities on a bank's balance sheet, are widely covered by the Financial Services Compensation Scheme, and so have a reduced incentive to scrutinise the bank's financial health. Secondly, banks often lend to other banks, who share an incentive to maximise their lending. Finally, regulators have been historically weak in maintaining rules on capital ratios, in the face of an industry equipped with large amounts of expertise in risk management. These factors combined to permit a situation in which bankers, acting on behalf of their shareholders, took (and were able to take) excessive risks.

Bankers' remuneration packages evolved to fulfil the needs of shareholders. Whilst much recent commentary has revolved around the incorrectness of the incentive structures this created, we argue that this was epiphenomenal, and ultimately caused by shareholder preferences. Shareholders of banks consistently voted for the remuneration packages that were put in front of them. Longer term remuneration packages (with "clawback" provisions for poor performance in later years) were not previously implemented, as shareholders were not prepared to accept the possible loss of personnel (and damage to shorter term returns) that would have resulted. Put another way, they shared both the investment horizon and the moral hazard of the bankers they employed – counter to the common perception that a principal-agency problem caused risk-taking distortions. We argue that a change in the ownership structure of a bank, and a consequent change of risk preference, would contribute to resolving the issue of banker remuneration, at the same time as resolving the problem of excessive risk-taking.

Externalities: Is the finance sector acting in the interests of industry?

For public welfare to be maximised, individuals and firms together need to be able to choose the mix of consumption, savings, borrowing and investment that is right for them over the long term.

The finance sector is the essential intermediary in this choice and in particular, it has a role in delivering the appropriate amount of capital for investment in future productivity.

There is room for debate about the correct level and form of that investment, but at present it appears that productive investment in the UK economy is running at levels below what is desired by citizens. That the investment level may be less than its long-term average is not in itself a problem - the choice made by consumers between current and deferred consumption may simply be different to the choice made in previous decades. But if the level of investment that consumers wish to make is not being achieved because of perverse incentives or behaviour within the finance sector, that is a problem with major long-term costs.

Are the financial markets, then, underproviding capital for productive investment? We argue that they are. The bias towards short-term returns and the nature of governance and reward structures in the finance sector make it harder for banks and institutions to extend long-term lending to corporates. In the short term this permits increased consumption but in the long term it will result in slower growth and reduced living standards. Business investment figures[1] suggest that UK investment has historically been, and remains, lower than that in most other countries.

[1] See for example http://www.bankofengland.co.uk/publications/speeches/2006/speech282.pdf, a speech by Sir John Gieve, for discussion and sources

Bounded rationality: Cognitive bounds of small (and large) shareholders

A model where a complex company is owned by a large and diverse group of shareholders will always suffer from problems of both information asymmetry and bounded cognition.

Rational ignorance

Bounded cognition of small shareholders (especially individuals, but also fund managers) creates problems of misaligned incentives or "rational ignorance". While better known as a problem in public choice economics, rational ignorance is of critical importance in corporate governance too.

The problem can be summarised as follows: the cost to a small shareholder of carrying out proper due diligence on a company, or even spending time to understand what it does, far outweighs any possible benefit. An individual who holds £2,000 of RBS shares might need to spend weeks educating herself to understand even a small part of the risks the company is taking, the likely future path of the credit markets it is involved in, and so on. But the cost of that investigation will be several thousand pounds in the investor's time alone. The maximum possible benefit she can achieve is to save their £2,000 investment by selling the shares before the company collapses, and so a simple cost-benefit analysis says that she should not bother investing the time to learn about the investment she is making.

The problem for larger investors is similar though at a different level. A fund manager holding £100 million of RBS shares can be expected to have carried out a certain level of research into the financial sector and the risks their shares are exposed to. But to carry out proper due diligence on their investment – of the kind that any venture capitalist would carry out on a startup they are investing in – would be so expensive for a company of RBS's size that it would outweigh any likely benefit. A likely cost running into millions of pounds could in theory be recouped if the share price were to move by a couple of percent, but the incentives within the fund management structure are likely to make this cost prohibitive. The fund manager is unlikely to have the freedom to spend his investors' money on such investigations, and thus they would have to cover it from their own fees – which will never be high enough to make it worthwhile, even if the value of the investment were to move by as much as 50%.

While large private equity investors and hedge funds might have sufficient incentives to carry out this kind of investigation, they have no incentive to make the results public. Therefore, the cost of acquiring information and the dilution of ownership create an undersupply of accurate analysis of large public companies.

Risk biases

A similar problem exists in the risk appetite of individual shareholders and the distorted risk incentives of institutional investors. Again the phenomena are distinct but related.

Individuals can be shown experimentally to behave irrationally with respect to risk. They consistently misprice risk by underestimating small probabilities. Individuals also tend to privilege the short term over the long, in a way that is inconsistent with their stated (and believed) preferences with regard to immediate consumption.

It can therefore be shown that the risk-taking behaviour of a large body of shareholders may not allocate capital in the most efficient way.

Risk-taking incentives for larger shareholders are distorted in a different way but with similar results. Reward structures for some institutional investors are asymmetric – they pay out on short-term profits but do not confiscate on loss; and those who charge a flat fee unrelated to performance are still incentivised to achieve high performance in the short term in order to attract more investment in the medium term. This phenomenon has been widely discussed with regard to the risk-taking behaviour of bank employees, but exists just the same among investors. This gives them an incentive to go along with high-risk behaviour within banks.

For these reasons – bounded rationality and risk mispricing – the financial sector is encouraged to concentrate more heavily on short-term investment than is economically optimal. As a result, there has historically been lower long-term investment than is optimal.

While these phenomena *prima facie* apply to any industry sector, there are two reasons why the finance sector is different.

The first is that other industries have financial professionals as a discipline on their activities. It is easier for a lender looking inwards to spot high risks – and withdraw the finance they offer – than for an insider to raise concerns within the finance industry and put their career at risk. To an extent this reduces risk-taking behaviour and short-term focus within other fields.

The second and more important is that excessive risk-taking and excessive focus on short-term returns in finance impose a negative externality on the rest of society. Part of the role of the finance sector in a well-functioning economy is to manage risks and transfer consumption, savings and investment between people and intertemporally. If its own operations do not rationally allocate risk this will have a powerful distorting effect on the rest of industry and on consumer behaviour. And if financial institutions fail, capital is immediately reallocated within the economy without notice, and into suboptimal places.

Similarly, a healthy finance sector creates positive externalities for the economy, by allowing capital to be invested where it brings the greatest returns, and allowing consumption to be smoothed by individuals over a long period.

These externalities argue for public involvement in or regulation of structures which are not supportive of stable risk allocation and investment. Naturally the nature of such public involvement needs to be limited, and should be specifically aimed at correcting identified distortions to minimise politicisation of the regulator's role.

Information asymmetry and compensation

Information is fundamental to markets. Information asymmetry is correspondingly a common reason for market failure. And information asymmetry in the market for financial sector compensation is a major problem.

Compensation is usually and rightly tied to individual and corporate performance. As a principle this is valuable but it relies on individual and corporate performance being accurately known to both parties. When performance is measured and announced by one party alone, there is an incentive for it to be overestimated - either wilfully or simply by selection and survivorship bias.

At the individual level this creates incentives for traders to misrepresent their performance or hide losses. However, internal governance within banks usually prevents this from becoming a systemic problem. At the corporate level the problem is harder to solve. Bank management have strong incentives to declare profits year after year - if necessary with an occasional $40 billion write-down which eliminates that year's bonuses but prepares the balance sheet for a surging recovery the year after.

The problem arises because it is bank management who control the published profit figures and benefit from the figures being presented in a certain way. Auditors have a role in controlling this, but the conflicts of interest in the audit sector are well-known, and in any case auditors will only ever have weak incentives to uncover problems. Similarly regulators do their best, but unless they observe every transaction in every market will not be able to provide a genuine market discipline.

Only lenders or investors have a real interest in identifying the true performance of a company, and as previously discussed the structure of the markets is such that those incentives are rarely correctly aligned.

Proposals

To resolve the problems of excessive risk taking and short-termism caused by limited liability shareholding and bounded rationality, we propose the following:

1. Attention to the ownership structures of banks and financial institutions
2. Incentives for long-term ownership of equity and loan instruments
3. Moves towards *meaningful* transparency – not just disclosure of data, but knowledge that can be effectively used by small shareholders and investors

We have not produced detailed policy proposals on these themes but we will outline the arguments in principle, and give some examples of potential policy actions.

Ownership structures

We propose two policy directions with regard to the ownership of financial institutions.

(1) The creation of more cooperative and savings banks, whose ownership arrangements are less susceptible to excessive risk-taking, would help to reduce the risk posed by limited liability ownership and reduce the systemic risk of the sector, by diversifying the business models employed. The argument for policy encouragement of different ownership structures is that the principal-agency problem combined with limits on farsightedness of economic agents distorts the market's normal preference for diversity; and that diversity, because it reduces systemic risk, creates a positive unpriced externality. This externality means that diversity is undersupplied by the market acting alone.

(2) The active employment of UK Financial Investments (in the banks which it owns a stake) to balance the possibility of excessive risk-taking of other shareholders in the future. Moreover, such activism could be used to address the issue of investor short-termism by tasking UKFI with championing longer term investments (possibly in areas of strategic importance to the government such as green technology). This could reasonably be achieved on a commercial basis, without the government directly interfering with the investment decisions to be made. If the government then chooses to reduce its shareholdings in the banks over time, this role could be replaced by an independent statutory body which could appoint board members from people with experience of industries outside the financial sector. This would serve both to break the cycle of excessive risk taking and help the UK develop specialisms in sectors that depend on stable long-term finance, where the UK economy has been historically weak.

In general the goal is not to have public ownership of the banking system but to have stable, committed owners who can help balance the interests of finance and the rest of the economy. Encouraging long-term shareholdings is one way to achieve this, as outlined in the next section.

Incentives for long-term ownership

In classical capital markets theory, there should be no difference between short-term and long-term ownership because the market fully discounts all expected long-term returns into current prices.

In reality, just as in the housing market, long-term ownership of financial assets creates incentives for investors to invest time in becoming well-informed about their holdings – an investment with positive externalities for the market as a whole.

In particular, long-term ownership substantially weakens the rational ignorance effect both by allowing time for investors to become well-informed, and by increasing the returns to such knowledge.

Thus, there is a strong case for incentivising long-term holding of both equities and bonds. This could be achieved with a taper-style capital gains tax or corporation tax benefit, or by adjustments to the stamp duty regime which would discourage short-term holdings. To minimise the distortive effects of such incentives, it may be plausible to direct them towards particular sizes of shareholding – on the theory that very small shareholders would be unlikely to carry out meaningful research anyway, and very large ones should already have sufficient incentive to do so.

Meaningful transparency

The obvious solution for many kinds of information asymmetry is greater transparency. By creating standards for disclosure, regulators can help restore a balance between market participants who have been disadvantaged by limited knowledge and those with privileged access to information. Openness – for example making public the terms of financial contracts that have been entered into by banks – also allows third parties to make more accurate judgment of counterparty risk and lets arbitrage impose additional market discipline.

However, two objections must be addressed.

The first is the competitive disadvantage suffered by regulated institutions in comparison to those elsewhere in the world. If UK banks, for example, are obliged to disclose more information about their contracts and positions than those elsewhere, the execution of some contracts will move offshore. One answer to this may be internationally coordinated regulation; another is mechanism design which gives either or both parties an incentive to disclose the terms of contracts. These issues do surface, of course, in the negotiations over international accounting standards and there are ways, albeit imperfect, to deal with it[2].

The second objection is that naive disclosure of information is not enough. Straightforward disclosure requirements in consumer finance do not always achieve the consumer benefit that they are intended to; consumers often have neither the time nor the understanding to correctly interpret large amounts of small print. Sophisticated investors are better placed to understand the contracts they are entering into, but still may not be good at evaluating large amounts of data about banks' other activities, even if such data were available. Indeed, this is one of the drivers behind innovation in finance and other markets – once a product has existed long enough to be well-understood by the majority of buyers, it starts to become commoditised and its suppliers suffer margin pressure. New products earn higher margins because buyers find it harder to compare them with alternatives.

Workable mechanisms for transparency

In the consumer space, a solution has been proposed by (among others) Sunstein and Thaler[3]: a central online portal containing information about credit card terms, allowing consumers to enter details about their situation and advising on the best option for their individual circumstances. This is a simple intervention which acts to improve the operation of competition in the interest of consumers.

Our proposal incorporates elements of this and other transparency mechanisms, with the goal of using the competitive pressures of the marketplace to complement regulatory actions.

The first step is to develop standardised contracts for the majority of financial products – notably those such as credit derivatives which are currently traded OTC. These contracts would be published by regulators, and the users of products would have an incentive to prefer the standard contracts

---

[2] An informative discussion of accounting standards can be found at http://www.voxeu.org/index.php?q=node/3610

[3] *Nudge: Improving decisions about health, wealth and happiness*, Richard H. Thaler and Cass R. Sunstein, Yale University Press, 2008

because they have familiar terms. Suppliers would come under market pressure to adopt these contracts even if they were not made mandatory.

Financial institutions would be required to publish aggregate figures showing the total size and parameters of these standard contracts they have entered into. This would give the market the ability to evaluate the overall risk profile of each institution. Currently the rating agencies play a role in the credit markets which is analogous to this but, to say the least, there is disagreement about the accuracy of their models. Publishing standard contract terms would allow market participants to develop and apply their own models to estimate risk.

Institutions would be free to enter into contracts on bespoke terms – although market pressure would likely put limits on this – but would be required to publish an accurate representation of the total scale and nature of such contracts. In particular, assets and liabilities which can currently be held off-balance-sheet would need to be brought onto balance sheet, acknowledged and fairly valued.

The "market for lemons" argument[4] indicates that enforced disclosure is in the interest of both buyers and sellers: without it, the market is stuck in an unproductive equilibrium where no seller can credibly disclose information voluntarily, due to free-rider problems. Thus, while no bank would voluntarily publish its contracts in the current equilibrium, there should not be undue political resistance to this proposal.

A detailed analysis of disclosure, regulation in general and their interactions with cognitive biases has been carried out by David Hirschleifer and Siew Hong Teoh[5].

Other incentives for research

The transparency argument interacts with the goal of long-term shareholding. As an alternative to relying on the research efforts of institutional shareholders (which may still be undersupplied if they do not share information), it would be possible to directly incentivise detailed research which is made available to the public. Of course, there is already an industry of equity analysts providing such research – if not to consumers, at least to investment bank clients – but they have had little credibility since the discovery of systematic biases in their advice, and allegations in a few cases of direct stock price manipulation, in the early 2000s. The ability to carry out research relies in itself on greater transparency, as it is questionable whether analysts can give meaningful advice without more access to detailed company data than they have at present.

It can be argued that the existence of hedge funds and private equity are the best incentives for research – because if correct, they are able to capture a large share of the benefits while correcting market prices. However, by its nature this research can only be funded by the losses of investors on the other side of the hedge funds' trades. This mechanism encourages prices to get out of line and become subject to sharp corrections rather than promoting stable and accurate valuations over the long term.

[4] The market for lemons: quality uncertainty and the market mechanism, George A. Akerlof, Quarterly Journal of Economics, 1970

[5] Hirschleifer and Teoh (unpublished), http://mpra.ub.uni-muenchen.de/14046/1/MPRA_paper_14046.pdf

Conclusions

There are inherent cognitive and structural problems in both the ownership and management of financial services firms which mitigate against efficient market outcomes.

An analysis of these problems shows why they lead to market failure and suggests policy actions which could move towards correcting them.

Current government ownership of banks through UK Financial Investments offers an opportunity to act carefully to achieve a public good, without creating a precedent for heavy-handed future intervention.

Similarly, regulation can be adjusted without restricting the freedom of market participants to design new products and freely enter into contracts. By increasing transparency and helping small shareholders and lenders improve their understanding of the risks taken by their agents, and of the balance between the short and long term, barriers to efficient market operation will be removed.