Black Swans and Financial Crises: Implications for Financial Regulation

John Quiggin
University of Queensland

1 October 2014

Abstract

How, if at all, is it possible for financial regulators to guard against unforeseen contingencies, given that, by definition, we do not when or how they will arise, or what their effects will be. In this paper, it will be argued that regulatory systems can be made more robust by imposing limits in financial innovation. The problem, then, is to distinguish between beneficial and harmful financial innovations. This is a difficult, but not impossible task

1 Introduction

The Global Financial Crisis of 2008, along with the subsequent sustained depression in the US and Europe, has changed the terms of debate over economic theory and financial regulation. One of the most striking features of the Crisis was the fact that it came as an almost complete surprise, most notably to those responsible for macroeconomic policy and financial regulation.

The years immediately before the crisis were notable for complacency. In 2008 a seminar on 'convergence in macroeconomics' concluded that the long-running debates between Keynesian and classical macroeconomists had been resolved, both in theoretical and in policy terms, broadly along the lines advocated by Milton Friedman.

Even more striking was the belief, generally held by central bankers that the business cycle had been tamed by a combination of improved monetary policy and financial deregulation. Termed the 'Great Moderation', this new era of economic stability was the subject of a rapidly growing academic literature. Quiggin (2010, Chapter 1) provides a detailed critique.

The (seemingly) sudden descent into a crisis that threatened the collapse of the entire global economy focused new attention on the role of unforeseen contingencies and incomplete awareness in generating instability in financial markets. The 'black swan' metaphor[1], popularized by Taleb's (2007) well-timed book of the same name, has become popular to describe such contingencies.

[1] Not, of course, original to Taleb

This discussion has, for the most part, taken place outside the mainstream of the economics profession and with little reference to formal decision theory. However, since the early 2000s, a substantial body of research has developed, dealing with the concept of unawareness[2].

Initially, the primary concerns of the literature were with issues of knowledge and belief, such as the question of whether an individual can know that there exist propositions of which they are unaware (Grant and Quiggin 2013a). More recently attention has been paid to problems of decision under bounded awareness (see for example, Grant, Kline and Quiggin 2012, Walker 2012) and to the implications for economic theory and policy.

Grant and Quiggin (2013b) consider differential awareness as a theoretical basis for the Precautionary Principle, frequently advocated as an approach to environmental regulation.

Quiggin and Siddiqi (2014) consider the relationship between bounded awareness and financial innovation. Quiggin and Siddiqi argue that incomplete and differential awareness plays a crucial role in the operation of financial markets. The interaction between awareness and arbitrage is crucial in explaining both beneficial and harmful financial innovations. Financial crises may be explained as the result of sudden changes in awareness, which may be exogenous or endogenous.

In the present paper, the ideas of Grant and Quiggin (2013) and Quiggin and Siddiqi (2014) are integrated to propose a precautionary approach to financial regulation.

2 Unawareness and financial markets

Research on financial markets and choice under uncertainty has been dominated by models derived from the work of Arrow and Debreu (1954) on general equilibrium theory and of Savage (1954) on decision theory. The central theoretical construct in these models is that of a state space Ω, a set describing all possible contingencies that may affect the outcomes of the decisionmaker.

Such unawareness may arise either because multiple contingencies are lumped together (coalescence) or because some contingencies are disregarded completely (restriction). Quiggin and Siddiqi (2014), drawing on earlier work by Grant and Quiggin (2103a) and Heifetz, Meier and Schipper (2006), address the problem by distinguishing between two kinds of unawareness.

On the one hand, as in Heifetz, Meier and Schipper (2006), agents may fail to distinguish between distinct states of nature, and may treat the them like a single state. We refer to this approach as coalescence. Under coalescence, the state space available to the agent is represented by $\tilde{\Omega} = \rho(\Omega)$ for some projection ρ.

Alternatively, agents may fail to consider some possibilities at all. Grant and Quiggin (2013) use the term 'restriction' to describe the associated state-space

[2] Schipper (2014) provides a bibliography

representation of bounded awareness, in which the state space considered by the agent is a proper subset $\tilde{\Omega} \subseteq \Omega$ of the full state space.

The contrast between unawareness as coalescence and unawareness as restriction may be illustrated with reference to a simple financial decision, considered by Quiggin (2014). Suppose an agent may choose to invest some of their wealth in financial assets and that there are two such assets, a bank account and a stock portfolio. Suppose further that there are two relevant propositions, s 'bank safe' (with $\neg s$ interpreted as 'bank fails') and u 'stock price up' (with $\neg u$ interpreted as 'stock price down'). These two propositions give rise to a state space Ω consisting of four states of nature

$$\Omega_{su} = \{s \wedge u, s \wedge \neg u, \neg s \wedge u, \neg s \wedge \neg u\}$$

A financially unsophisticated agent might be unaware of the existence of stock markets, and would therefore display pure unawareness with respect to the proposition u. The agent would therefore consider only the state space $\Omega_s = \{s, \neg s\}$. From the perspective of a more aware agent with access to Ω_{su}, the states $s \wedge u$ and $s \wedge \neg u$ are coalesced into a single state s and similarly for $\neg s$. With multiple agents who must impute awareness to each other, this process gives rise to a lattice of state spaces, as discussed by Heifetz, Meier and Schipper (2006).

Alternatively, an agent might be aware of both banks and stock markets, but might not consider the possibility of bank failure. Such an agent would consider the state space is $\Omega_u^s = \{s \wedge u, s \wedge \neg u\}$, a restriction of Ω_{su}.

In the financial sector example, both possibilities are relevant. Some agents may be unaware of the existence of particular financial securities and may not distinguish different events relevant to the payoffs of those securities. For example, an agent who is unaware of the existence of stock markets (or, more realistically, has no idea of how they might go about investing in stocks) need not be aware of events like 'stock markets go up'. Alternatively, investors may regard distinct securities as being effectively equivalent, when their potential consequences are very different.

For example, in the leadup to the Global Financial Crisis, a variety of mortgage-backed securities were rated AAA; that is, were regarded as effectively replicating Treasury bonds. As it turned out, the set of possible consequences, including default in the context of a deep recession, were very different.[3]

Thus, we have a dual relationship. Restriction of awareness regarding the state space implies coalescence of elements of the space of acts and restriction of the set of consequences. Coalescence of awareness requires restriction of the set of acts or coalescence of the set of consequences.

2.1 Coalescence

Consider first the case when unawareness takes the form of coalesence of the state space. An agent can only consider assets with returns that are measurable

[3] I am indebted to a referee for suggesting this point. Some related issues are discussed in Quiggin (2013).

with respect to her partition of the state space. Conversely, an asset can only exist if there is at least one agent for whom the return vector is measurable.

In these circumstances, it is easy to describe situations where arbitrage opportunities are not exploited. An example will illustrate. Suppose the asset market includes a bond which returns 1 with probability 1, and the equity of three companies competing to develop a new product. One of the three companies will succeed, returning 1 while the other two will fail.

Observe that a portfolio consisting of all three equity assets replicates the bond. Hence, normalizing the price of the bond to 1, and assuming symmetry, the no-arbitrage condition requires the price of the equity assets to be equal to 1/3.

Suppose, however, that all agents in the economy have coarse awareness and more particularly, that each agent is aware only of one company, and can therefore hold only the bond and that company's equity. Risk aversion implies that the price of the equity assets will be less than 1/3. Hence, the equilibrium price vector will violate the no-arbitrage condition. This example illustrates:

Proposition 1 *In an economy where unawareness takes the form of coalescence, arbitrages can persist in equilibrium*

2.2 Restriction

Suppose that the spanning and no-arbitrage conditions are satisfied, and consider an agent with restricted awareness, who is unaware of some state ω. By the spanning condition, it is possible to design a portfolio $a^{-}(\omega)$ with payoff -1 in state ω and 0 in all other states, and this portfolio will have a price of $-\rho(\omega)$.

Now consider any portfolio a with a non-negative payoff in all states except possibly ω, and a price which may be normalized to 1 without loss of generality. By combining a with $\frac{1}{\rho(\omega)}$ units of $a^{-}(\omega)$ we obtain a portfolio with a cost of zero and a non-negative payoff in every state other than ω. From the viewpoint of an individual unaware of ω, this is an arbitrage. Hence, we obtain

Proposition 2 *Suppose there exists a complete spanning set of securities and no arbitrages. Then an agent who is unaware of some states will incorrectly perceive arbitrage opportunities*

Corollary 3 *If all agents display unawareness of some state, the equilibrium outcome will violate the no-arbitrage condition*

2.3 Unawareness and financial innovation

For several decades after the deregulation of the 1970s, financial innovation was seen as unambiguously beneficial. As with technological innovation in general, it was agreed that innovation might disrupt established ways of doing things and impose losses on those slow to adapt. However, again as with technological innovation, it was assumed that the benefits to society as a whole would outweigh these losses.

The standard theory of financial markets presents them, simply as exchanges for insurance contracts. Individuals endowed with vectors of state-contingent income trade among themselves to achieve a socially optimal allocation of risk, aligned with the aggregate social endowment. Financial innovation is a beneficent process involving the creation of new opportunities for exchange.

This description captures an important part of reality but it is, obviously, partial. The reality of financial markets includes frenetic bubbles and busts, in which previously unknown or obscure assets rise dramatically in value, before becoming worthless. Sometimes, the effects of these bubbles and busts are confined to participants in the markets in question. At other times, as in the present prolonged depression, they bring ruin and destitution to millions of people who have never knowingly traded a financial instrument more complex than a savings deposit.

Quiggin and Siddiqi (2014) argue that incomplete and differential awareness plays a crucial role in the operation of financial markets. The interaction between awareness and arbitrage is crucial in explaining both beneficial and harmful financial innovations. Financial crises may be explained as the result of sudden changes in awareness, which may be exogenous or endogenous.

The central observation of Quiggin and Siddiqi (2014) is that the equilibrium outcome when agents have bounded awareness violates the no-arbitrage condition. This fact creates profit opportunities for agents with greater awareness. Such agents can exploit arbitrage opportunities, at least up to the point where other agents become more aware, or when prices shift so as to restore the no-arbitrage condition.

Arbitrage opportunities may be exploited through trade in existing assets. However, more comprehensive exploitation will typically involve financial innovation, that is, the creation of new derivative assets.

In the example given for coarse awareness, an agent who can exploit the arbitrage opportunities, for example by creating a derivative asset that uses equity to replicate the bond, will bring asset prices closer to their equilibrium values, producing a net improvement in social welfare. Such innovations illustrate the potential benefits of financial markets in achieving a socially superior allocation of risk.

On the other hand, all agents are boundedly rational. Innovations may arise because agents with restricted awareness incorrectly perceive opportunities for arbitrage. In this case, the market equilibrium may move far away from the first-best. Such innovations are undesirable.

In a sense, undesirable innovations are self-correcting. They create opportunities for more aware agents to profit at the expense of those holding the overvalued assets. However, there is no reason, under conditions of bounded rationality, to suppose that the process of correction will be smooth. Moreover, both because of regulation and because of the nature of financial institutions, sharp corrections in financial markets frequently create negative externalities. Banks that make bad investment decisions may either be bailed out or may default on unrelated obligations, leading to systemic crises. Hence, there is no reason to suppose that the problem of undesirable financial innovations will be resolved smoothly by financial market processes.

Rather, the analysis presented above provides an outline for the way in which bubbles and busts may be generated in financial markets. First, a new financial instrument is introduced, which appears to exploit existing arbitrage opportunities but in fact (or in addition) creates new ones.

The development of the new instrument creates incentive for individuals to increase their awareness in ways that allow them to discover and exploit arbitrages. These agents will be keen to sell the new instrument short, while less aware agents will wish to buy it. Thus, the volume of trade can grow rapidly, creating a bubble.

The bubble bursts when enough people become aware of the previously unconsidered possibility. Once short positions are large enough, it will be in the interests of those holding them to make others aware of this fact.

3 The precautionary principle

Grant and Quiggin show how the precautionary principle, often advocated in the context of environmental decisionmaking, may be justified on the basis of bounded awareness. In the context of environmental policy, 'innovations' such as new projects and processes may have unforeseen adverse consequences. Hence, it may be ecologically rational, in the sense of Goldstein and Gigerenzer to reject such innovations, even when an a benefit-cost analysis (which necessarily excludes unforeseen consequences) is favorable. Grant and Quiggin argue for a modified version of the precautionary principle in which innovations with unforeseeable consequences may be approved if the potential losses from adverse unforeseen events can be bounded appropriately.

The analysis of Grant and Quiggin may be adapted to the problem of regulating financial innovation.

Represent the problem as an extensive-form game with Nature, denoted Γ, using notation based on that of Osborne and Rubinstein (1994) and Halpern and Rego (2006b). Nature moves first, and makes two independent choices. The first, from the set $\{b, \neg b\}$, is whether or not the system behaves in the way foreseen by decisionmakers. Nature chooses b with probability π_b. The second, from the set $\{s, \neg s\}$, is whether or not the financial system suffers a shock, such as the bursting of a housing bubble. Nature chooses s with probability π_s. The 'black swan' (b, s) case with probability $\pi_b \pi_s$ arises when there is a shock and the system behaves in a way unforeseen by decisionmakers.

After nature's move, but before being informed about it, the decisionmaker makes two choices. First, player 1 (the decisionmaker) chooses from the set $\{t, \neg t\}$ where t denotes undertaking a 'stress test' at cost C to ensure that all reasonably possible consequences have been considered. More precisely, the stress test will determine which of $\{b, \neg b\}$ Nature has chosen and therefore whether a 'Black Swan' shock is possible. If the study reveals a 'Black Swan' possibility, the outcomes associated this possibility may be incorporated into a cost-benefit analysis.

Second, player 1 chooses from the set $\{i, \neg i\}$ where i is a financial innovation. This innovation will yield positive returns in the absence of a shock, a small loss if there is a shock and the system behaves in the way foreseen by decisionmakers, and a large loss in the 'black swan' event.

Notationally, we represent a history h as a sequence of actions. So, for example, the sequence of actions 'Nature chooses Black Swan and player 1 chooses no stress test' is denoted by $\langle b, s, \neg t, i\rangle$. A complete history, such as $z = \langle b, s, \neg t, i\rangle$. is referred to as terminal and associated with a payoff $v^1(z)$ for player 1. The payoff for player 1 is the sum of three terms: a benefit B, received if the project is approved, a cost C incurred if the study is undertaken, and damage H, L or 0, determined by Nature's choice and incurred if the project is approved.

More precisely, the payoffs are

$$\begin{aligned}
v\langle *, *, \neg t, \neg i\rangle &= 0 \\
v\langle *, *, t, \neg i\rangle &= -C \\
v\langle *, \neg s, \neg t, i\rangle &= B \\
v\langle *, \neg s, t, i\rangle &= B - C \\
v\langle b, s, \neg t, i\rangle &= -L \\
v\langle b, s, t, i\rangle &= -L - C \\
v\langle \neg b, s, \neg t, i\rangle &= -H \\
v\langle \neg b, s, \neg t, i\rangle &= -H - C
\end{aligned}$$

The information sets $I^1(h)$ available to player 1 at history h (that is, the histories considered possible by player 1 when the true history is h and 1 is to play) are: $\{\langle *, *\rangle\}$, after Nature's move; $\{\langle *, *, \neg t\rangle\}$ after a decision not to test, $\{\langle b, *, t\rangle\}$ after a test revealing b; and $\{\langle \neg b, *, t\rangle\}$.

A strategy β^1 for 1 is a rule prescribing an action a at each information set. A belief system π^1 assigns to each information set I^1 a probability distribution over its elements, and a pair (β, π) constitutes an assessment.

Consider the strategic choices for a decision-maker fully aware of the problem structure. A strategy β^1 for 1 is a rule prescribing an action a at each information set.

It easy to see that the undominated strategies are:
β_1 no test, no innovation
β_2 no test, innovation
β_3 test, innovation at $\{\langle b, *, t\rangle\}$, no innovation at $\{\langle \neg b, *, t\rangle\}$

$$\begin{aligned}
V_1 &= 0 \\
V_2 &= (1 - \pi_s) B - \pi_b \pi_s L - (1 - \pi_b) \pi_s H \\
V_3 &= \pi_b (1 - \pi_s) B - \pi_b \pi_s L - C
\end{aligned}$$

We will assume that the parameters are such that $V_3 > 0 > V_2$.

Incomplete awareness is modelled by the assumption that the decision-maker, initially represents the problem by the subjective game structure Γ' which does not include the possibility that nature will choose $\neg b$ at move 1. It follows that the decision-maker is unaware of the high damage consequence H.

Thus, the perceived information set is $\{\langle b, *\rangle\}$ after Nature's move and $\{\langle b, *, \neg t\rangle\}$ after a decision not to test. If the decisionmaker undertakes the test, the information sets are $\{\langle b, *, t\rangle\}$ after a test revealing b; and $\{\langle \neg b, *, t\rangle\}$ after a test revealing $\neg b$.

An incompletely aware decisionmaker who disregards the possibility of the decision $\neg b$ perceives the payoffs as

$$\begin{aligned} V_1^* &= 0 \\ V_2^* &= (1-\pi_s)\,B - \pi_s L \\ V_3^* &= (1-\pi_s)\,B - \pi_s L - C \end{aligned}$$

With the probabilities as before, we have

$$V_2^* > V_3^* > V_3 > V_1^* = 0$$

Hence, if the decisionmaker disregards his incomplete awareness, he will choose the sub-optimal strategy β_2 thereby risking the large loss H.

Now consider a boundedly rational, but sophisticated decisionmaker. Such a decisionmaker will be conscious of her own limited awareness and of the possibility that there exist hazards she has not considered that would cause her, *ex post*, to regret her decisions.

Suppose the decisionmaker, with subjective model Γ', understands that this model is incomplete and that decisions based on the model Γ' may have unforeseen consequences. What, if anything can be done about this.

Grant and Quiggin (2012) propose an alternative approach in which the deductive analysis of decision theory is constrained by heuristics derived from inductive reasoning, expressed in syntactic terms rather than the semantics of a fully-specified state space.Given past experience of decisions that have turned out badly because of unconsidered possibilities, the decisionmaker may regard the proposition 'there may exist hazards I have not considered' as being justified by induction.

Grant and Quiggin propose two heuristic interpretations of the Precautionary Principle in the context of environmental decisionmaking based on the concepts of surprises and secure strategies.

From the perspective available to someone (say, a modeller, or the decisionmaker following a study) who has access to the full model Γ, the most important history omitted from Γ' is $\langle \neg b, s, \neg t, i\rangle$. We will refer to a history of this kind, involving an unconsidered action by Nature as an 'surprise', and say that the strategy β_2 in which $\neg t$ and i are always chosen is 'subject to surprises'.

Surprises may be favorable or unfavorable. In problems where the precautionary principle is considered, inductive judgement will normally justify the proposition that some strategies are subject to unfavorable surprises. To make this notion more precise, we assume that the restricted game Γ' includes a 'secure' strategy which is not subject to surprises. In the example, the natural choice is the strategy β_0. Since it does not involve a change to the established set of financial instruments, this strategy involves risks that are well understood.

We refer to the expected payoff from the 'secure' strategy as the 'security level' for the restricted game, and adopt the convention (applicable to the example) that the security level is zero.[4]

We may define a strategy β as 'subject to unfavorable surprises' if the proposition 'there exist unconsidered moves by nature, against which the payoff from β is less than zero' is inductively justified on the basis of past experience.

We consider the following Strong Form of the Precautionary Principle:

> '*In games with a secure strategy, reject strategies that are subject to unfavorable surprises.*'

As Grant and Quiggin observe, the definition of surprises is very strong. In the present setting the Strong Form of the Precautionary Principle implies that strategies involving stress tests would also be subject to unfavorable surprises relative to the secure strategy of rejecting all innovations.

This suggests the need for a Modified Form of the Precautionary Principle in which strategies subject to unfavorable surprises may be adopted provided that an appropriate bound on the losses can be determined. In the present context, the Modified Form of the Precautionary principle is satisfied if the maximum loss from an unfavorable surprise is less than the expected benefit in the case where there is no surprise, given by $V^*(\beta)$. The idea is that, even with bounded rationality, the decisionmaker understands the problem well enough that the proposition 'It is less likely than not that $\hat{\beta}$ will suffer an unfavorable surprise' is inductively justified.[5]

Strategy β_3 (test, then innovate if b), incurs a loss of C in the 'surprise' event β Hence, β_3 satisfies the Modified Form of the Precautionary Principle provided that

$$(1-\pi_s)B - \pi_s L - C > C$$

or

$$(1-\pi_s)B - \pi_s L - 2C > 0$$

4 Implications for financial regulation

As Grant and Quiggin observe, Precautionary Principle typically advance in the context of contested public decisions, involving interactions between parties with different beliefs and interests. A common process of this kind is one in which a regulated financial institution introduces a new kind of security. In

[4] In this example, the payoff from the 'secure' strategy is risk-free. In general, however, secure strategies may be subject to risk arising from moves by Nature or other players that is explicitly represented in the restricted game Γ', and to which probabilities have been assigned.

[5] In some sense, this is implicit in the normal usage of the term 'surprise'

this context, the adoption of the precautionary principle by the regulator would entail rejection of all innovations where, given the regulator's level of awareness, the strategy of allowing the innovation was subject to strongly unfavorable surprises.

In this context, the precautionary principle may be seen as placing the burden of proof on the proponents of innovation. If the proponent can produce sufficient information, and therefore a sufficient increase in awareness, to satisfy the regulator that the innovation is not subject to strongly unfavorable surprises, then the project can proceed. Otherwise it is rejected. This is a reversal of the position, characteristic of 'light-handed regulation' , in which there is no requirement to seek regulatory approval before introducing innovations. Rather, the innovations take place first, and regulations are adjusted as problems emerge.

4.1 Too big to fail or too interconnected to fail ?

Critical discussion of the failures of financial regulation before and during the Global Financial Crisis have focused on the idea that some institutions are 'too big to fail' in the sense that their collapse would create too much economic damage to be permitted. Markose (2009) and Quiggin (2012) argue, on the contrary that the problem is that institutions are 'too interconnected to fail' in the sense that the failure of one institutions may impose unsustainable losses on its counterparties.

The distinction between 'too big to fail' and 'too interconnected to fail' may be understood in terms of unawareness. Consider a 'stand-alone' financial institution such as a building society, which takes deposits from savers and lends them out to borrowers against security such as houses. Such an institution may fail as a result of bad luck or imprudent management. The resulting failure will imposes losses on shareholders[6], employees, creditors and depositors (or, in the presence of deposit insurance, the insurer). The 'value at risk' approach, popular among regulators in the leadup to the GFC, provides a straightforward method of estimating the potential loss, given by the maximum failure rate of loans multiplied by the proportion of the loan amount not recovered in foreclosure. Value at risk may be large or small, but, given appropriate prudential supervision, does not typically involve 'black swan' events.

Since equity markets are widely diversified, and subject to large fluctuations, the loss of shareholder equity in even the largest corporate failure is unlikely to pose any systemic risk to either the financial system or the macro-economy. Hence, the 'value at risk' approach is generally adequate to deal with failures of stand-alone institutions. All that is necessary is to ensure that the combination of shareholder equity and deposit protection is sufficient to cover the value at risk. Deposit insurance schemes must also take account of the risk that a sharp downturn in an asset class (eg housing) may cause the simultaneous failure of

[6] At one time building societies were mostly co-operatives so the shareholders and depositors were the same people. However, for the purposes of this exercise, it is useufl to focus on the case of an investor-owned institution.

a number of financial institutions, but this is a problem separate from 'too big to fail'

Now consider, by contrast, the case of a financial institution which is involved in large-scale transactions with other institutions. Notable examples from the recent past include the hedge fund, Long Term Capital Management, which failed in 1998 and the investment bank Lehmann Brothers which failed in 2008[7] These institutions were lightly regulated, on the basis that their owners and depositors were sophisticated high-wealth individuals who could manage the risks involved. However, when they failed it became evident that their counterparties in various large-scale transactions (the banks which had lent to LTCM and the insurer AIG, which was exposed to credit default swaps) would also be at risk of failure, as in turn would be there counterparties.

In the terms of the model presented above, the risks associated with a highly interconnected financial system are hard to foresee and almost impossible to bound in advance. It might well have been possible to let LTCM and Lehmann Brothers fail, as the regulatory system was committed to do before the possible consequences became apparent. Equally, the failure of these institutions might have destroyed the global financial system entirely.

In retrospect, it was clear that the light-handed system of financial regulation, which permitted banks to expose themselves to large losses from unregulated health funds and untested financial instruments was unsound. However, the corrections made so far have the quality of shutting the stable door after the horse has bolted. While they are probably adequate to prevent an exact repetition of the crisis of 2008, they retain a presumption in favour of financial innovation that is inconsistent with the Precautionary Principle.

4.2 Policy response: narrow banking

Regulated financial institutions, such as banks, want the security that goes with (implicit or explicit) government backing but are unwilling to forgo opportunities for profitable investment, even where these increase risk. Examples of these practices abound. The result is a chronic problem of regulatory gaming.

Regulated banks have set up or acquired unregulated subsidiaries, then called for rescue when the subsidiaries have run into trouble. Banks have marketed a wide range of financial assets, with varying degrees of depositor protection, through the same branches that accept guaranteed deposits. The effect is to blur the meaning of the government's guarantee. Worst of all, banks have used their 'too big to fail' status to become investors on their own account, turning business risks into systemic risks. The result, evident in recent months, is the classic combination of capitalised profits and socialised losses.

[7] An earlier example, influential in my own thinking on this topic was the Finance Corporation of Australia, a subsidiary of the Bank of Adelaide and one of a number of finace companies that failed in the 1970s. While the others were allowed to fail, the FCA collapse endangered the parent bank, leading the Reserve Bank to step in and organize a rescue by ANZ.

Attempts to address these problems through regulations such as the Basel I and II capital adequacy requirements have failed completely at a global level. Although Australia's regulatory system has held up fairly well so far, there are large losses still to come, and serious vulnerabilities from our dependence on overseas finance.

The most plausible solution is a system of 'narrow banking' where regulated and guaranteed banks are confined to a specific set of low risk activities with well-established accounting rules. Financial innovation in this sector should be tightly restricted to ensure that new financial assets do not produce the kind of risk transfer to the public seen in the current crisis.

The narrow bank sector should be strictly separated from other financial markets, such as equity markets, and governments should make a clear commitment that non-bank financial institutions will never be allowed to become 'too big to fail', or rescued if they run into difficulties.

5 Concluding comments

Six years after the Global Financial Crisis, the problem of designing financial regulation to be robust to unforeseen shocks has barely been addressed, let alone resolved. The nascent theory of decision under unawareness provides one possible approach to this problem.

6 References

Arrow, K. and Debreu, G. (1954), 'Existence of an equilibrium for a competitive economy', *Econometrica*, 22(3), 265-90.

Grant, S., Kline, J. and Quiggin, J. (2012), 'Differential awareness, ambiguity, and incomplete contracts: A model of contractual disputes', *Journal of Economic Behavior and Organization*, 82(2-3), 494-504.

Grant, S. and Quiggin, J. (2013a), 'Inductive reasoning about unawareness', *Economic Theory*, 54(3), 717-755.

Grant, S. and Quiggin, J. (2013b), 'Bounded awareness, heuristics and the Precautionary Principle', *Journal of Economic Behavior & Organization*, 93, 17-31.

Heifetz, A., Meier, M. and Schipper, B. (2006), 'Interactive Unawareness', *Journal of Economic Theory*, 130(1), 78–94.

Quiggin, J. (2010), *Zombie Economics: How Dead Ideas Still Walk Among Us*, Princeton University Press, Princeton.

Quiggin, J. and Siddiqi, H. (2014), Differential Awareness, Financial Innovation and Financial Crises, Working Paper, University of Queensland.

Savage, L.J. (1954) *Foundations of Statistics*, Wiley, New York.

Schipper, B. (2014), 'The Unawareness Bibliography', http://www.econ.ucdavis.edu/faculty/schipper/unaw

Taleb, N. (2007) *The Black Swan: The Impact of the Highly Improbable*, Random House, New York.

von Neumann, J. and Morgenstern, O. (1944) *Theory of Games and Economic Behavior*, Princeton University Press, Princeton.
Walker, O. (2012). Unawareness with "possible" possible worlds, mimeo., London School of Economics.