Risk Management: Making God Laugh

In May 2004, preparing for a conference on terrorism, I typed the single word "risk" into Google. I got 40 million hits. For purposes of comparison I typed in "God" (60 million hits) and "sex" (80 million). I repeated the exercise for this article in December 2005. God got 149 million hits, sex 222 million, and risk 673 million. Risk, in a short period of time, has overtaken its principal competitors by a wide margin.

Figure 1

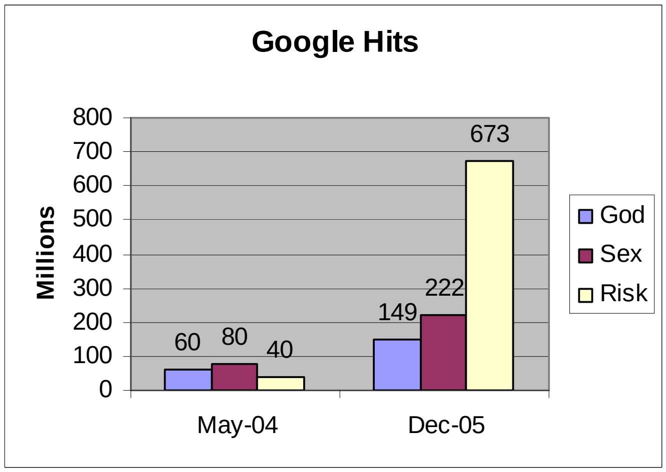

What is going on? How might one interpret these statistics? They provide a crude measure of what I suspect is a real, global, phenomenon – or at least a phenomenon that can be found in that part of the globe connected to the Internet. This qualification may be important; Google, which can be used to chart the growth of interest in risk, may also be feeding the concerns embodied in the word.

Risk is a word that refers to the possibility of something nasty happening in the future, and the future exists only in our imaginations. The Internet provides instant access to vast numbers of nasty possibilities to worry about, and also large numbers of people offering to hold your hand while you do. "Risk management" yields over 50 million Google hits, most of them oblivious to the wisdom of the Woody Allen joke about how to make God laugh – tell him your plans.

One need not sample all the websites concerned with risk to discover many unnecessary, and often acrimonious, arguments. People are using the same word "risk", but understanding different things by it, and shouting past each other. Figure 2 presents a typology that helps to clear away some of the unnecessary acrimony.

Figure 2

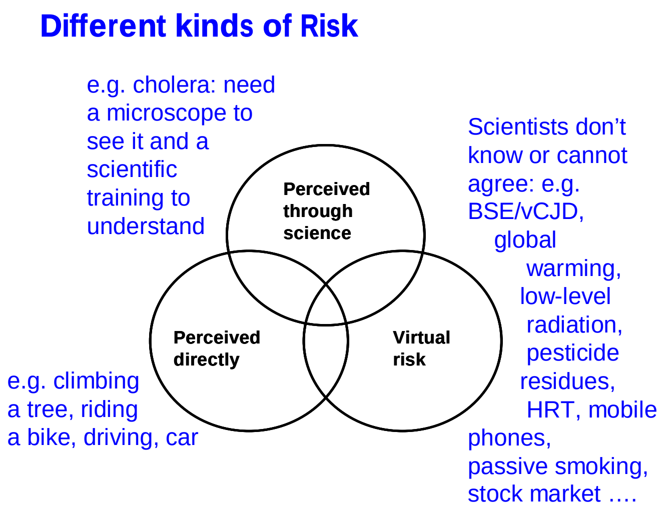

Directly perceptible risks we manage ourselves - if I am late for dinner, and see my bus approaching on the far side of the road, I will risk shorter gaps in the traffic in order to cross the road to catch it. I do not undertake a formal probabilistic risk assessment before I cross the road. Such risks are managed by the application of *judgement* – a still mysterious mix of instinct, intuition and experience. Thus far it has seen me safely across the road.

There are other risks that cannot be seen by the naked eye. Cholera, for example, can only be seen with the help of a microscope and by someone with a scientific training who knows what he is looking at. Science and technology have impressive deity-defying risk-management achievements to their credit. From molecular biology, through medicine and engineering, to statistics and epidemiology they have contributed to enormous increases in average life expectancy – and consequent, yet-to-be resolved, pensions crises. God may yet laugh.

But there is a third, much larger and more challenging category – *virtual risk*. Here we encounter the longest-running and most acrimonious debates about risk. If science cannot settle an issue, everyone feels liberated to argue from their pre-established beliefs, convictions or superstitions. This category presents particular difficulties for those proffering their services as risk managers. To manage something is to direct or control it. Another name for virtual risk is *uncertainty* - imaginable possibilities for which we have insufficient evidence to attach meaningful odds. And beyond uncertainty in this circle lie Donald Rumsfeld's unknown unknowns – now known, with the benefit of hindsight, to have thrown his plans for Iraq into chaos.

A recent paper in *Science*[1] trains a scientific "microscope" – in the form of functional magnetic resonance imaging – on to this problem; fMRI can detect which parts of the brain are excited by different sorts of problems. The amygdala and orbitofrontal

[1] Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making, Ming Hsu, Meghana Bhatt, Ralph Adolphus, Daniel Tranel, Colin Camerer, *Science*, 2005, vol 310.

cortex are of particular interest. They are implicated in the integration of emotional and cognitive inputs. The paper demonstrates – with pictures of the scans – that the greater the uncertainty attaching to a task, the more these parts of the brain light up on the scans – of normal people. People with lesions on these parts of the brain are apparently much less troubled by uncertainty and ambiguity.

Such people, the paper observes, are "behaviourally abnormal", but behave, ironically, in a way consistent with "the logic of subjective expected utility theory." The authors note that "in subjective expected utility theory, the probabilities of outcomes should influence choices, whereas confidence about those probabilities should not". They conclude that there is a part of the brain that can cope with odds, as calculated by a bookmaker, but works hard when confronted with uncertainy

An example they give involves two decks of red and blue cards. Deck 1 has 10 red and 10 blue cards (called the "risky" deck, because the odds are known). In Deck 2 (called the "ambiguous" deck) the numbers of each colour are not known. A bet on a colour pays $10 if right and $0 if wrong – yielding an average expected return of $5. Alternatively a subject can decline to bet for a sure gain of $3: so the "rational" risk manager should bet. In experiments with undamaged subjects people revealed a preference for placing a bet – red or blue – against the risky deck. This the authors call a paradox because, despite the uncertainty attaching to the composition of the ambiguous deck, a person schooled in Decision Theory would recognize that in both cases the chances of winning are the same as those attaching to the flipping of a coin. They conclude that their findings support "the hypothesis that ambiguity lowers the anticipated reward of decisions" and fosters risk averse behaviour.

How might this knowledge about how undamaged brains work in the face of uncertainty assist a risk manager familiar with expected utility theory? In the example discussed above the authors refer to their card experiment as "pitting pure risk (where probabilities are known with certainty) against pure ambiguity." But of course it did no such thing. In both cases the probabilities were not only knowable, but identical. Their experiment works only on people who have not figured this out. The problem that they construct to represent uncertainty still belongs in the scientific circle of Figure 2.

Consider a simple model of risk decision-making – the Risk Thermostat.

Figure 3

Risk Thermostat with Perceptual Filters

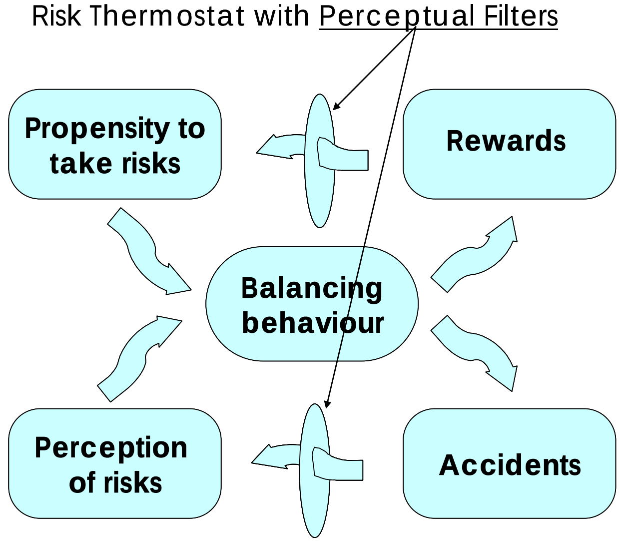

The model postulates that we all have some propensity to take risks, the setting of the thermostat in the top left-hand corner. The setting can vary widely from person to person depending on personality and circumstance. This propensity leads to risk taking behaviour, which leads, by definition, to accidents; to take a risk is to do something that carries with it a probability of an adverse outcome – you hope for heads, but it might turn up tails. Through having accidents and surviving them and learning from them, or seeing them on television, or being warned by mother, we acquire a perception of what is safe or dangerous. The model postulates that when perception of risks and propensity to take risks are not in balance the imbalance leads to behaviour that seeks to restore the balance. Why do we take risks? There are rewards for taking risks, and the magnitude of the reward influences propensity.

In Figure 3 the model has been fitted with perceptual filters. The filters are composed of instinct, intuition and experience, and moulded by culture. The less conclusive the science relating to any particular risk, the more influential becomes the influence of these filters. The institutional context in which risk decisions are made is also important. When I cross the road I perform the process described by Figure 3 inside my own head. I am the judge of the magnitude of the reward for catching the bus and the risk of being hit by a passing car.

When risk management becomes institutionalised there are strong pressures to replace judgement with calculation – with formal, probabilistic risk assessment. Standard risk assessment forms require the assessor/manager to identify risks and the "associated control measures". Effective control requires a firm grasp on the thing being controlled. Where unambiguous knowledge of cause and effect is not available, knowledge of the odds is the next best thing.

Institutional pressures commonly produce *bottom-loop bias*. Institutions confronted with ambiguity commonly respond by manifesting a complete disregard for the rewards of risk taking. The job specification of institutional risk managers usually makes them responsible for reducing accidents. Frequently they are enjoined not to have their judgement about what is safe or dangerous compromised by, or corrupted by, contemplation of the rewards of risk taking. The mantra of Britain's Health and

Safety Executive, the Britain's foremost risk manager is "Reducing Risks, Protecting People". It is invoked so often by institutional risk managers that it is sometimes reduced to R2P2 to save ink.

Less commonly, especially in certain types of financial institution, one encounters incentive structures that produce *top-loop bias*. If, in a good year, your Christmas bonus is large enough to retire on, and if an "accident" that loses your clients a large amount of money results only in your needing to find another job, you are likely to be a risk-seeking trader. But, as with your bottom-loop counterpart, you will seek to rationalize your behaviour with some form of probabilistic risk assessment.

We must be careful to distinguish virtual risks (true uncertainty and unknown unkowns) from risks about which scientists and decision theorists can offer helpful advice. When dealing with risks found in the virtual circle of Figure 2 we are thrown back, as in the directly perceptible circle, on *judgement* – a combination of instinct, intuition, and experience. Risk managers given the task of controlling or directing uncertainty are understandably risk averse. How, if you do not know the odds, or the rules of the game, or the name of the game, or whether your game is part of someone else's much bigger game (unknown unknowns) can you be expected to control or direct it … whatever *it* might turn out to be?

A common response to this dilemma is denial - to behave like the brain-damaged subjects in the experiments described by Hsu *et al*, to behave as though you know the odds. This can be achieved by producing a model (a guess about the structure of the problem you might be dealing with) and assigning probabilities to key variables in the model (more guesses). This can produce the sense of satisfaction associated with purposeful endeavour. But when it leads to the neglect or denial of significant variables not captured by the model, it can lead disastrously astray.

I offer two high-profile examples of the risks run if judgement is substituted by calculation. The first is compellingly documented by Roger Lowenstein in *When Genius Failed* , the story of the spectacular fall, in September 1998, of Long Term Capital Management, a fall that came close to bringing down the global financial markets. The principal "geniuses" in this story were Robert Merton and Myron Scholes who shared a Nobel Prize for Economics in 1997 for their discovery of "a new method to determine the value of derivatives". So long as the assumptions embodied in their model held, so long as the phenomena they were modeling could be confined within the scientific circle of Figure 2, their genius trumped all competitors, and produced astonishing profits. But their vanity, arrogance and early success deceived them into believing that they had a formula for managing uncertainty.

The second, in the news as this is written, is a petition that has attracted large numbers of signatories, objecting to the award of this year's Nobel Prize for Economics to Robert Aumann and Thomas Schelling for enhancing "our understanding of conflict and cooperation through game-theory analysis". The "understanding" of one of these prize winners (Schelling) the petition protests, was the direct inspiration for the US strategy in Vietnam of indiscriminant bombing that resulted in 2 million civilian deaths, and ended in ignominious defeat. And the expertise of Aumann, the petition complains, is currently being deployed as an argument against Israel's withdrawal from Gaza – a zone of conflict in which the combatants label each other "terrorist" while claiming God for their side.

Chess is a game with two players and precise rules governing the moves they might make. The Deeper Blue computer, which achieved a narrow win over Kasparov, was capable of analysing 200 million moves a second. For games with contested rules involving multiple players - whether terrorists or stock market traders - the computer, for the foreseeable future, will lag a long way behind. If we attempt to apply the methods appropriate to the circle illuminated by science, to questions that they cannot answer, we behave like the drunk searching for his keys, not where he dropped them, but under the lamppost because that is where it is light. If God is amused by human presumption, that surely will make Him laugh.

Figure 4 Risk management: where are the keys?