PUBLIC PERCEPTION OF RISK

J. Richard Eiser
Centre for Research in Social Attitudes
Department of Psychology
University of Sheffield
Sheffield S10 2TP
UK

Email: j.r.eiser@shef.ac.uk
Tel. 0114-222-6622.
Fax. 0114-276-6515.

REPORT PREPARED FOR
FORESIGHT
OFFICE OF SCIENCE AND TECHNOLOGY

JULY 2004

This report was commissioned by the Foresight programme of the Office of Science and Technology. However, its findings are independent of government and do not constitute government policy.

Executive Summary

This review considers public perception of risk from a perspective based primarily on psychological theories of attitudes, decision-making, learning and social influence.

Part I presents the theoretical framework. Part II applies this to three specific contexts, or 'case studies', relevant to particular Foresight projects: the Flood and Coastal Defence Project, the Cyber Trust and Crime Prevention project and the Brain Science, Addiction and Drugs project.

The central themes are how people make judgements about possible future events under conditions of uncertainty, and the decisions they make to achieve desired outcomes and avoid undesired outcomes.

Risk is a feature of all human action which has effects that are more or less uncertain and yield some kinds of benefits or costs. So perception of risk involves implicit or explicit judgements of the likelihood or uncertainty, and the desirability or undesirability, of such effects.

Such judgements are made both by decision-makers and by those affected by their actions (e.g. the public). Their judgements and decisions are based on expectancies derived from their own experience and from information communicated to them by others. These expectancies guide 'approach' and 'avoidance' behaviour. If people expect an activity or product to be beneficial on balance, they will tend to approach - engage in a behaviour, or use a product. If they expect it to be damaging or costly, they will generally prefer avoidance - reject it or leave it alone.

People's expectancies may be confirmed or contradicted by the consequences of their decisions, but there is a fundamental asymmetry between approach and avoidance behaviour. Approach behaviour can produce informative feedback whereas avoidance behaviour does not. If you avoid something you believe to be dangerous, you will not discover whether your fears were justified.

This can lead to a bias in the direction of 'false-alarm' or risk-averse responding, which will be even more marked where the perceived benefits of accepting the risky activity or product are unclear.

However, expectancies about consequences can also be over-optimistic. This is frequently the case where actions (or products) are associated with both costs and benefits, but the costs are delayed (e.g. unhealthy behaviour), or are inconsistently experienced (e.g. accidents following dangerous driving, or being caught speeding). Although people's judgements about the riskiness of activities or products can be updated through experience, such experience is typically selective and incomplete, and will tend to be interpreted as consistent with people's prior beliefs.

Attitudes and behaviour (risk perceptions and decisions) can also be influenced by other people. Because our friends tend to share our attitudes, and we are more likely to accept advice from our friends, such influence can reinforce existing attitudes.

Trust can involve reliance on others, both as controllers of risks and as informants about the extent of risk. In both cases, trust can depend on implicit estimates of the others' competence, values or partiality and honesty.

If 'experts' are seen as having a vested interest (for example if they are in the pay of industry, or committed to defending the standpoint of a political party or pressure group), this may undermine trust.

Decision-making quality may be undermined by pressures for premature closure and failure to consider alternative hypotheses. Within organisations, tolerance of minority opinions and encouragement of innovation is valuable. Communicators found to have withheld information lose trust.

PART I.
Uncertainty and Choice

Risk is regarded as a central concern in many policy debates. How can risks be identified and measured? How can they be managed? When should they be accepted or rejected? And most commonly, how are they likely to be interpreted or perceived by different people? These questions arise in areas as diverse as health and lifestyle, hazardous industries, pensions and investments, transport, climate change and environmental protection, and in many other less obvious contexts. An internet search engine recently took 0.18 seconds to find "about 30,400,000" uses of the word risk. This tells us that the term is pervasive. But it would be amazing if it were being used in the same way in each of these 30 million instances.

It is tempting simply to acknowledge that this is a huge field, and admit that it is impossible to generalise about it. Why should the decisions of air traffic controllers, investment managers and cigarette smokers be thought to have anything in common just because they can all be said to involve some kind of 'risk'? The safe approach, on this approach, is to treat each instance separately and focus on the concrete details of each case.

But there is an alternative, which is to trust our intuitions and to accept, at least as a working hypothesis, that different kinds of risk have something in common, and to attempt to analyse what it might be. If it succeeds, this approach will allow us to go beyond description towards a more complete theoretical explanation, and perhaps even permit us to make predictions of relevance to areas that have so far been less researched

The aim is to identify unifying principles or processes, if any can be found. It is not sufficient to document what we know about different kinds of risks and see

what they might have in common. If the results of our web search tell us anything, it is that the sample of instances is too huge to sample systematically. Instead, we shall start from the other end, outlining a few very general principles, before moving on to see how they might be applied in some specific cases.

My central theme is that everything that is important about risk arises from actual or perceived uncertainty. There can be uncertainty about both the likelihood of events and their value or desirability. These two elements are basic to the definition of risk, and guide our actions in any context where risk is involved. Indeed, it is only because we need to act under conditions of uncertainty that the concept of risk is of interest. If we felt there was nothing we could ever do to affect what might happen to us, we would have no decisions to take and there would be no point in worrying about the likelihood or value of future events. However, most of the time life isn't like that. We have choices to make, and these choices can have consequences for ourselves and others. It is because these consequences are uncertain, and may leave us better or worse off, that we talk about risk.

The terms 'risk-taking' and 'risk-aversion' (or 'avoidance') are used to describe different kinds of choices under uncertainty. The phrase 'perception of risk' refers to how we anticipate the outcomes of choices made by ourselves or others. The 'events' that we appraise in terms of their likelihood and scale are very much ones that are either brought about, exacerbated or ameliorated by human decisions and human actions. In short, the perception of risk is an example of the interpretation of uncertain information, and risk-taking and avoidance are kinds of actions chosen under conditions of uncertainty. Understanding such actions is part of the more general question of how cognitive systems – both individual and social – handle and respond to uncertainty in a changing world.

Much has been said about the differences between the opinions held about any given risk by the general public on the one hand and experts on the other. Such differences can be real enough, but they are a red herring from the point of view of an understanding of underlying processes. Despite this, it is the views of the public about risk, rather than those of experts, that have been and continue to be of most concern to policy-makers. Why is this?

One practical consideration is that the desirability of winning public acceptance can be a major constraint on how, and even whether, particular policies can be enacted. This is evident in objections of local residents and environmental groups to new industrial or infrastructure developments (e.g. wind farms, airport runways, waste incinerators, mobile phone masts, or flood defence schemes). Public consultation can be costly and time-consuming, even where the eventual outcome is favourable from the perspective of the policy-makers. The perception of risk (particularly to health, but also to a fragile natural or cultural environment) has a special place in such debates in that it introduces an ethical dimension. It is one thing to acknowledge that any development will produce economic losers as well as winners, and that therefore some deal needs to be done. It is quite another matter to maintain

that an increase in the incidence of childhood cancer, or the destruction of a rare wildlife habitat, is a price worth paying for economic progress. The discourse of risk is very powerful, whatever lies behind it.

In the same way, risk-management decisions may be influenced by public perceptions of risk in ways that may distort priorities away from actual risk reductions. Policy-makers may feel the need to be seen to be doing something about particular risks, even where the risks are relatively small and the actions undertaken are more visible than they are effective. A possible example is in the area of crime prevention. Recent UK surveys suggest that, whereas the incidence of many forms of crime is falling, public fear of crime appears to be increasing (Hale, 1996). This can slant policies towards initiatives such as more CCTV in shopping areas, and 'more Bobbies on the beat'. These measures may make people feel safer, which is itself a benefit, but may also lead to less actual crime reduction since crime may be displaced to other areas, and police officers may be diverted from more effective forms of crime prevention.

The other side of the story is that the general public can sometimes appear frustratingly complacent about the seriousness of other kinds of risks, and so be resistant to policies and actions that could lead to risk reduction. Prime examples come from many aspects of health behaviour such as poor diet choices, lack of exercise, smoking, alcohol and substance abuse, and sun (and sun-bed) exposure. Lower speed restrictions may be resisted by motorists' organisations. Measures to reduce carbon emissions may be unpopular with many individuals, corporations and governments. Mostly such complacency is expressed in terms of a down-playing of the risk rather than as a flat denial. For example, the great majority of smokers (at least in the UK and US) accept that smoking is somewhat damaging to health, but may underestimate the extent of the risk.

These examples show that policy-makers need a better understanding of how the public perceives risk in general, and specific risks in particular. But it is dangerous to assume that there is something fundamentally different about the way in which risks are perceived by 'the public' on the one hand and by policy-makers or 'experts' on the other, with the expert view being taken as a benchmark of objectivity with which the former is contrasted. People all take on different roles at different times: of ordinary citizens, of local residents, of 'experts' with special experience, or of decision-makers needing to balance benefits and costs for ourselves and for those dependent on us.

There is increasing recognition, therefore, of the need to move away from the traditional research agenda focused on questions such as how public perceptions differ from expert opinions, the potential reasons for such differences, and what can be done about them. This research has stressed better communication to reduce the discrepancy, typically by shifting public perceptions towards those of the experts. Not only does this approach pose the wrong questions, it makes many researchers wary of how their answers may be interpreted. A bleakly cynical view could be that much risk perception research has been a thinly disguised exercise in social control, directed

towards manipulating public opinion so that it is brought into line with the wishes of the powers-that-be. Mindful of being tarred with that brush, some researchers have instead affected an attitude of agnostic neutrality, according to which all views of risk should be given equal credence as subjective representations of alternative realities.

I do not believe we need to let the pendulum swing back that far. Experts and the public often come to different conclusions about the seriousness of any particular risk. But although they may use similar processes to reach their conclusions, it would be wrong to suggest that there is nothing to choose between the conclusions which these groups reach in terms of their correspondence with objective evidence. It is absurd to suggest that experts are omniscient, and never make mistakes. But arguing over whether the risk assessments of experts are 'better' than those of non-experts is a red herring. Either this must be trivially true (since making good assessments is part of what it means to be an expert) or else the argument is about something different – the validity of the criteria according to which we attribute 'expertise'.

Furthermore, there is a vital distinction between prediction and prescription. Experts should be better at predicting what will happen than non-experts, but recommending what action should be taken on the basis of such predictions necessarily rests on some view of what outcomes are most desirable. Patients consult a doctor both for diagnosis (prediction) and for advice or treatment (prescription). Patients and doctors are engaged in a formalised social encounter governed by cultural norms and expectations on both sides. Crucially, patients are entitled (and expected) to trust their doctor to share their own goals for what the consultation should help achieve. They want to get better and their doctor wants them to get better too. But if the different parties to an encounter have different values and priorities, the prescriptions of the experts may be rejected, however convincing their predictions. An example might be a choice between economic growth and environmental conservation

Uncertainty and chance in science

Uncertainty is a key concept in thinking about risk. But is uncertainty a state of mind – a reflection of our own incomplete knowledge - or is it intrinsic to the subjects about which we seek knowledge? A range of attitudes to this question can be found in different scientific disciplines, and within disciplines over time, but possibly all sciences are more tolerant of some notion of chance than popular conceptions of science might suggest.

Statisticians believe in chance, and seek to quantify it. To this end, they distinguish between two main types of uncertainty that can apply to predictions concerning any future event. The first ('epistemic', from the Greek word for understanding) arises from lack of knowledge, while the second ('aleatory', from the Latin word for dice) reflects 'pure chance', or the randomness of any sampling from a population. To distinguish between them we need to know about the distribution we are considering. Suppose, for

instance, that an epidemiological study suggests that one in every ten smokers is likely to suffer from a particular disease within a given time-frame. (This is roughly the proportion who get lung cancer.) Does this mean that each and every smoker has the same one-in-ten 'chance' of getting the disease? Or does it mean that the population of smokers contains a minority who are almost bound to get the disease if they smoke and a majority who are comparatively immune? As we discover more about genetic risk factors for cancer and other diseases, the second alternative can no longer be dismissed.

Within the physical sciences, there is a long tradition of deterministic world-views. The 18th Century French mathematician Pierre Simon de Laplace argued – before computers were invented - that an intellect vast enough to analyse data on the positions and movements of all physical objects could come up with a 'single formula' so that nothing would be uncertain and the future would be as evident as the past. For Laplace, everything in the universe obeys deterministic laws, if only we could find them, so the only kind of uncertainty is epistemic. However, with the advent of quantum physics, we now know that, at the quantum level, physical objects and events demand alternative descriptions and that the very act of measuring something can alter what is being measured. Furthermore, even deterministic processes are not always certain or predictable. They can produce essentially unpredictable or 'chaotic' outcomes (Stewart, 1989). Hence, the argument goes, however much we learn about the causes of things, the future behaviour of systems and organisms will remain unpredictable.

Physical scientists, nonetheless, tend to be universalists by outlook. One of the great projects of contemporary theoretical physics is the search for a complete unified 'Theory of Everything'. Behavioural and social scientists, by contrast, are a rather mixed bag, and have more modest ambitions, at least in public. Their disciplines are concerned primarily with human behaviour, and there is a broad consensus that this is a pretty complex business. Any behaviour we try to predict is likely to be influenced by a huge number of factors, only some of which we can hope to control or hold constant. In predicting human behaviour, we never reach the point where we can say "We've controlled for all possible situational factors, and all personality differences among our participants, so all the remaining variation reflects pure chance, rather than the influence of some additional, as yet unidentified, cause". The distinction between epistemic and aleatory uncertainty has fewer practical implications for the social sciences. If we do not know the extent of our own ignorance it is impossible to apportion our uncertainty between ignorance and chance.

The practical impossibility of settling such issues empirically does not prevent social and behavioural scientists from holding heated debates about the nature of chance in human behaviour. Do we have free choice, or is it an illusion? Is all our behaviour merely determined by our previous conditioning? What is the relative importance of environmental constraints, family background and inherited characteristics on human behaviour and achievement? Even if they are difficult to resolve, these debates are more

than philosophical games. They have implications for how we attribute responsibility to individuals for their actions. From a methodological point of view, they affect how we should view the validity of different kinds of data, particularly individuals' self-reports concerning their motives, and the bases for their preferences and choices. Increasingly, it is being acknowledged that much human behaviour is the consequence of rather rapid 'automatic' decisions which are made without a great deal of reflection and are relatively inaccessible to introspection (Bargh, 1997; Fazio, 2001).

Uncertainty and chance in everyday life

A fair amount of psychological research has looked at how people seek causal explanations for particular events, and especially how much they allocate responsibility to factors such as personal agency, situational constraints and chance (Heider, 1958; Ross, 1977). Apportioning blame for an accident, for instance, turns greatly on how much it is believed that a particular person or organisation did things that caused the accident to occur, or failed to do things to prevent it, and whether the consequences of such actions could be foreseen (Eiser, Reicher & Podpadec, 1995a).

Other work has identified differences in the extent to which individuals believe that they are in control of their own destiny, rather than being at the mercy of influences beyond their control. The belief that one is helpless to prevent negative consequences through one's own behaviour has been claimed to be associated with clinical depression (Alloy, Abramson & Francis, 1999; Dweck & Goetz; 1978). Related research suggests that many individuals feel threatened by the notion of a world in which bad things can happen just by chance. An unattractive instance of this is the phenomenon known as 'victim derogation'. We are worried (so the story goes) about the idea that mishaps and victimisation can happen to anyone at random, since this means that we are just as much at risk as anyone else. We find it more comfortable to believe that victims' lack of prudence (or even moral worth) meant that they, at least partly, brought the mishap on themselves (Lerner, 1980).

This suggests that the notion of chance is a problematic one. Both in science and in everyday social interaction, appeals to chance appear to be explanations of last resort. Even in situations where the outcomes are meant to be random, such as lotteries, and some forms of gambling, many people seem attracted to superstitious rituals or 'systems' that provide them with illusions of control (Langer, 1975). In business and government, this need for prediction and control is manifested in formal systems of risk analysis and risk management, and in the design of procedures that are flexible enough to deal with the unexpected (Strategy Unit, 2002). But risk analysis and risk management are a part of everyday decision-making and are as vital to individual survival as they are to the success of organisations.

But our uncertainty about the future is based in reality. The world we inhabit is hugely complex. This is true of the principles that govern our physical environment, and of the way in which we interact with each other and with the other intelligent creatures and complex organisms with which we have co-

evolved. Yet somehow we have not merely evolved, but flourished. Our endowment of cognitive, emotional and behavioural capacities has enabled us to find our way around a less than perfectly predictable world and succeed in getting enough of what we want.

Cognitive heuristics

Processes can be adaptive for survival and success without being error-free. In fact, a highly influential programme of research in cognitive psychology over the last 30 years has provided ample evidence of how human decision-making can be vulnerable to various forms of 'bias' and 'error'. Since much of this evidence takes the form of inconsistent processing of statistical information, this area of work would appear to have direct relevance to the field of risk perception. The key notion is that, when faced with problems requiring statistical inference or judgment under uncertainty, people tend to rely upon so-called 'cognitive heuristics', in other words rules of thumb or short cuts (Tversky & Kahneman, 1974). These often provide good approximations to a correct answer but under some conditions can lead to significant errors.

The early thrust of this research was to identify some of these heuristics and illustrate some of the errors which can ensue. For example, the 'availability heuristic' refers to the tendency to assume that events are more probable if they are easier to recall or imagine (more 'available' or 'accessible' to memory). Generally speaking, more frequently experienced events will be easier to recall or imagine, but the converse need not always hold. One suggestion is that vivid reporting of uncommon events such as some kinds of crime can make people think that they are common. Another well-researched heuristic is termed 'representativeness'. This involves thinking that a characteristic shared by most of a particular category of objects or people is more typical than it is. Just because most secretarial staff in an office happen to be female, this doesn't mean that most female office staff are secretaries.

In order to explain why we are vulnerable to the influence of heuristics, researchers such as Nisbett and Ross (1980) made reference to our relatively limited capacity for consciously processing several pieces of information at the same time. This means that we need to be selective in what we attend to and remember. The advantage of heuristics was that they provide fairly good answers at a fraction of the cognitive cost (Hogarth, 1981). More recently, the emphasis has shifted to showing how these heuristics are actually well-suited to helping us make decisions in more naturalistic contexts (Gilovich, Miller & Kahneman, 2002). Gigerenzer and Hofrage (1995) observe that people appear much less prone to errors of statistical reasoning when provided with information in the form of frequencies rather than probabilities. They argue that this corresponds more closely to how we build up expectations over time.

We are also recognising that heuristics may allow people to take decisions involving emotional and motivational processes that can provide a guide to action as well as being quicker and easier than fully-developed thought processes. Indeed, the steady removal of the barriers between the study of

emotion on the one hand and cognition on the other has been one of the most important theoretical developments in psychology over the last 25 years (e.g. Bower, 1981; Damasio, 1994). Slovic et al. (2002) use the term 'the affect heuristic' for the tendency for people to rely on their own immediate emotional reactions and evaluative associations when forming preferences and judgements, including perceptions of risks. This idea builds on previous evidence that our first impressions that something is 'good' or 'bad' are formed extremely quickly, often without our being able to report exactly what has prompted such feelings (Zajonc, 1980). According to Slovic, an implication for risk perception is that people may find it difficult to estimate benefits and costs separately. If they 'feel good' about an activity or issue, they will judge it to have more benefits and fewer costs; if they 'feel bad' about it, they will judge it to have fewer benefits and to be more costly and dangerous. Research on attitude change has shown that people frequently accept or reject persuasive messages on the basis of 'heuristic' as opposed to 'systematic' processing (Chaiken, 1980) of the information provided, essentially by focusing on less informative or 'peripheral' cues and associations (Petty & Cacioppo, 1986). People are more likely to rely on evaluative associations rather than on a more careful calculation of consequences when they feel that an issue is less personally relevant to them (Petty & Cacioppo, 1990), or where they lack the motivation or capacity to process the information more systematically (Fazio, 1990).

Prospect theory

But despite these advances in our understanding, research on cognitive heuristics has primarily involved judgements about the likelihood of possible events rather than their value or desirability. Allegations of 'error' and 'irrationality' in such judgements are based on comparisons with normative principles of statistical reasoning. However, since risk involves both likelihood and value, what happens when some form of value is introduced into the decisions with which research participants are presented? Here again we have a normative definition of 'rationality', that offered by classical economic theory. It proposes that, if different possible outcomes are associated with measurable values and probabilities of occurrence, then a 'rational' decision strategy is to calculate the 'expected value' for each outcome my multiplying up the probability of the outcome and its value, which could be positive or negative. The option associated with outcomes of higher expected value should then be preferred, and individuals should be indifferent between options of the same expected value.

Is this what happens? Not always. Research inspired by Kahneman and Tversky's (1979) Prospect Theory demonstrates that expressed preferences can be strongly influenced by contextual features which are supposedly irrelevant to the expected value of different outcomes.

Kahneman and Tversky (1979) employ a version of a method known as the 'standard gamble'. This involves asking individuals to choose between options associated with the same expected value, but differing in their level of certainty. In non-technical terms, this is like choosing between a 'bird in the

hand' and 'two in the bush'. If the chance of catching the two birds in the bush is 50:50, the two options should have the same expected value – one bird - and hence be equally attractive. The data (like the proverb) suggest otherwise. Faced with analogous choices, people tend to prefer the sure thing (the bird-in-hand), unless the chances of winning with the gamble (catching the two birds-in-the-bush) are substantially improved. Kahneman and Tversky refer to this as a tendency for individuals to be 'risk-averse for gains'. However, this effect appears to be reversed if individuals are faced with a choice between losses. When faced between accepting a sure loss of $10 and a 10% probability of losing $100, but a 90% chance of losing nothing, participants tend to prefer to take the gamble. In other words, they are 'risk-seeking for losses'.

These effects depend to a great extent on how the problem is described ('framed' in the research jargon) to the participants. Individuals' preferences in a particular situation can be shifted between risk-aversion and risk-seeking by the way in which it is presented. A choice could be expressed in terms of gains, such as a choice between public health policies which may result in different numbers of lives being saved, or in terms of losses, such as the number of people who may die (Fischhoff, 1983; Kahneman & Tversky, 1984; Tversky & Kahneman, 1981).

Framing also influences judgements of probability. Tversky and Koehler (1994) argue that possible events come to be seen as more probable if they are 'unpacked' into separate components making up that class. For instance, Fischhoff, Slovic and Lichtenstein (1978) presented participants (including experienced car mechanics) with 'fault trees' of the possible causes of a car's failure to start. When categories of causes were specified in more detail (e.g. "battery charge insufficient" was unpacked into "faulty ground connections", "terminals loose or corroded", "battery weak") they were assigned a higher likelihood. Johnson et al. (1993) presented participants with hypothetical health insurance decisions, and similarly found a greater willingness to pay for protection for hospitalisation for "any disease or accident" than for "any reason". Tversky and Koehler (1994) suggest that added details can make possibilities easier to imagine whereas catch-all categories (e.g. "other engine problems" or "all other problems") are more difficult to think about.

Caution is needed in extrapolating from such experimental findings to real life. Participants in these experiments make judgements about risk, but they are rarely at risk in any of these situations. The choices made are predominantly hypothetical, involving imaginary rather than actual dilemmas. Are we necessarily risk-seeking for losses? It depends. It is easy to imagine gamblers, or investment managers in the mould of Nick Leeson (the former 'rogue trader' at the Singapore office of Baring's Bank), who get more and more into debt because of a reluctance to accept a moderate loss when they believe that, if their luck would just turn one more time, they could redeem themselves.

On the other hand, we are prepared (up to a point) to undertake various forms of defensive expenditure, on insurance premiums, burglar alarms, and lock-up

garages, all of which involve the acceptance of a definite cost to offset the consequences of a larger, but uncertain, loss. And we do this often in the knowledge that the companies who sell us such protection are making a healthy profit, i.e that the expected values are in their favour. Are we necessarily risk-averse for gains? Again, it depends. Gambling and lotteries are popular pursuits, and gamblers would not need to advise each other to 'quit while you're ahead' if they were never tempted to do otherwise. And of course, all gamblers know that bookmakers, casinos and lotteries collect more in stakes than they pay out.

However, there is an even more fundamental difficulty with using these kinds of experiments as models for everyday decision-making. We know in advance the values and probabilities of the outcomes associated with each different outcome. We know how much we stand to win if we stake £100 on a horse starting at odds of eight to one. In everyday life we have to ask: "What are the chances of this or that happening?" " How good or bad would it be?" or "How can we find out?"

Experiential learning

The relevance of experimental research on cognitive heuristics and decision-making under uncertainty to real-life risk perception and risk-taking, therefore, is less well established than might appear at first sight. We need to know more about how we learn about our environment and develop expectancies and predictions from experience. This experience is necessarily selective. Could learning processes account for the development of tendencies such as risk-seeking and risk-aversion, and possibly the use of cognitive heuristics too? This is an intriguing possibility, but as yet there has been comparatively little research that attempts to bring the fields of learning and decision-making into contact with each other.

An important exception is work on experiential learning. The essential feature of experiential (or reinforcement) learning, in other words learning from experience, is that the learner has to do something in order to gain information. For simplicity, let's assume that we have two choices when faced with an unfamiliar object or situation: to approach or avoid. We may have some initial expectations that lead us to anticipate pain or pleasure (say), and these could guide our decision, or we could simply guess. If we choose to approach the object, our action will allow us to find out more about what the object is really like, and whether the consequences of interacting with it are painful or pleasurable. As a result of this experience, we can update or correct our initial expectations, and these in their turn will guide our future behaviour. If we derive pleasure from our experience of the object, our expectations will be more positive the next time we encounter it, or something similar, and we shall be more likely to approach it again (approach will be 'reinforced'). If we experience pain, our expectations will be more negative and we shall be less likely to approach it a second time.

The implication is that over time, our expectations about different objects should fall broadly into line with their actual value. But such learning only

occurs if we choose to approach objects in the first place. If we avoid objects indiscriminately, we can never find out whether they are truly good or bad.

The principle applies much more widely than simply to human learning. In studies of reinforcement learning in animals, a related concept is termed the exploration/exploitation trade-off (Sutton & Barto, 1998). Exploration is equivalent to approach, for instance visiting a new arm of a maze (in a laboratory experiment) or a new tree or field (in a natural environment) to see if it contains an attractive food source. Exploitation involves the animal using its existing knowledge about the location of food sources to approach (or dangers to avoid) so as to maximise its immediate outcomes. The trade-off between these two forms of motivation rests on the facts that exploration involves risk but exploitation is only adaptive in the short term, since known food sources may become depleted.

This process can also be illustrated in artificial systems. March (1996) used computer simulation to argue that risk aversion for gains could be a direct product of learning experience, rather than any kind of personality trait or higher-order conscious process. His simulations were based around the 'standard gamble' dilemmas of the kind used by Kahneman and Tversky (1979). A learning system was presented with two classes of objects. The 'sure thing' option always produced a moderate reward, while the 'risky' option sometimes produced a much larger reward, but more often produced nothing. The reward magnitudes and probabilities were arranged so that the expected value associated with the two classes of objects was the same. The learning system had to choose whether to 'approach' or 'avoid' these objects on the basis of its expectations developed over time. The learning algorithms used meant that these expectations were modified by experience of the value of individual objects, if and only if they were approached.

March's findings showed strong evidence for risk aversion for gains. The argument goes something like this. Whenever an object of the 'sure thing' class is presented, it is consistently associated with a good outcome, and so approach is strongly reinforced. Before long, all 'sure thing' objects will be approached. By contrast, most 'risky' objects will produce no good outcome or reward at all, and so the tendency to approach these will be weakened. As a consequence, 'sure thing' objects will be more likely to be approached and 'risky' objects avoided.

"Ah, but what about the fact that the two classes of objects have the same expected value?" an economist might ask. "Shouldn't the extra value of the occasional large rewards in the 'risky' class compensate for their infrequency?" True, but this only applies if both classes are fully sampled. Once the system starts differentially to approach the 'sure thing' class and avoid the 'risky' class, the chances are that it will never sample enough of the 'risky' class to discover that it contains a few high rewards. Unlike human participants in standard gamble experiments (or lottery players, perhaps), the learning system doesn't 'know' that there's a jackpot out there, unless and until it finds it. This raises an interesting question about the human experimental evidence. Why should participants act as though they have to

experience the value of objects for themselves, even when they have been told about the presence of a less likely jackpot?

A similar idea underlies some of our own research, involving both human experiments and simulation (Fazio, Eiser & Shook, 2004; Eiser et al., 2003). Instead of the standard gamble, our learning situation involves a 'virtual world', presented in the form of a computer game. Participants are told to imagine that they are in a world containing different kinds of 'beans' (these are shown varying in shape and numbers of speckles). Some beans are good and provide energy, others are bad and make one sick. The participants' task is to find out which beans to eat and which to avoid, something they can only do if they first sample the beans. If they fail to find enough good beans, they will 'die'. Thus, learning involves risk, but failure to learn is also fatal. The basic finding is that participants are quite good at identifying the bad beans and avoiding them, but they only identify a proportion of the good beans, and continue to avoid some good beans as though they were bad. The same pattern is observed in the computer simulations, when the learning system is trained by an algorithm that requires an approach analogous to approach (or 'eating') for feedback (i.e. reinforcement) to occur.

The importance of such findings for an understanding of risk perception and risk-taking is that:

- learners are (initially) uncertain about the values associated with different objects or actions
- they acquire expectancies about these values
- these expectancies guide their decisions to approach or avoid particular objects
- approaching objects can have either good or bad consequences: it involves risk
- experiencing good or bad consequences leads to modification of expectancies, i.e. to learning
- hence, learning involves risk
- exploration guided by such expectancies leads typically to an incomplete and selective sampling of the available information.

Costs and benefits

The idea that we process information selectively is not new. It underlies work on cognitive heuristics and goes a fair way towards explaining several observed biases and inaccuracies in human judgement. But is accuracy the only benchmark against which we should evaluate the effectiveness of an information-processing strategy?

Other things being equal, any form of reasoning or information processing that leads to more accurate predictions will be preferable to one that leads to inaccurate predictions. But other things are not always equal and accuracy is not the only goal. Some errors will be more costly than others and some correct choices will be more beneficial than others.

The balance between anticipated costs and benefits is central to risk-taking (Lopes, 1987). Risks don't just happen. They are brought about by human activity, sometimes unwittingly but frequently by someone somewhere along the line estimating, rightly or wrongly, that the benefits associated with a course of action outweigh the likely costs. Asbestos, we now know, can be fatal if its fibres are inhaled. But this is not the 'fault' of asbestos. Asbestos constitutes a risk to human health because of how it has been used in building construction, with the commendable aim at the time of safeguarding people and property against fire. Even natural disasters such as earthquakes, floods and volcanic eruptions only constitute risks because of decisions taken by people about where to live, how strongly to build their houses, or whether to invest in flood protection. And human activity can make such 'natural' events worse than they otherwise might have been, for instance through climate change.

A framework for considering the costs and benefits of different decisions derives from a classic theory of visual perception known as Signal Detection Theory (SDT, Swets, 1973). The basic problem this theory addresses is that of describing the 'discrimination performance' of somebody faced with the task of identifying whether or not a piece of information is evidence of a 'signal' or merely 'noise'. For instance, how does a radar operator tell the difference between a blip on a screen due to an approaching aircraft and one due to atmospheric disturbance? How reliably can a safety inspector (or an automatic device such as a smoke detector) distinguish danger from safety? How well can a doctor diagnose a particular condition from a set of clinical symptoms?

SDT distinguishes two parameters of performance: sensitivity or discrimination ability (the proportion of correct responses, in other words accuracy) and criterion or response bias (the tendency to give a response in one direction, for example to say the 'signal' is present, or that the patient has the disease, regardless of the actual facts of the matter). Depending where the criterion is set, some ambiguous pieces of information may be overinterpreted as a 'signal', and others wrongly discounted as 'noise'. Figure 1 illustrates the effects of adoption of a risky criterion, where the signs have to be quite strong to be interpreted as dangerous, and a cautious criterion, where the situation is only declared safe if possible signs of danger are absent or very weak.

Figure 1. Illustration of a signal detection problem.

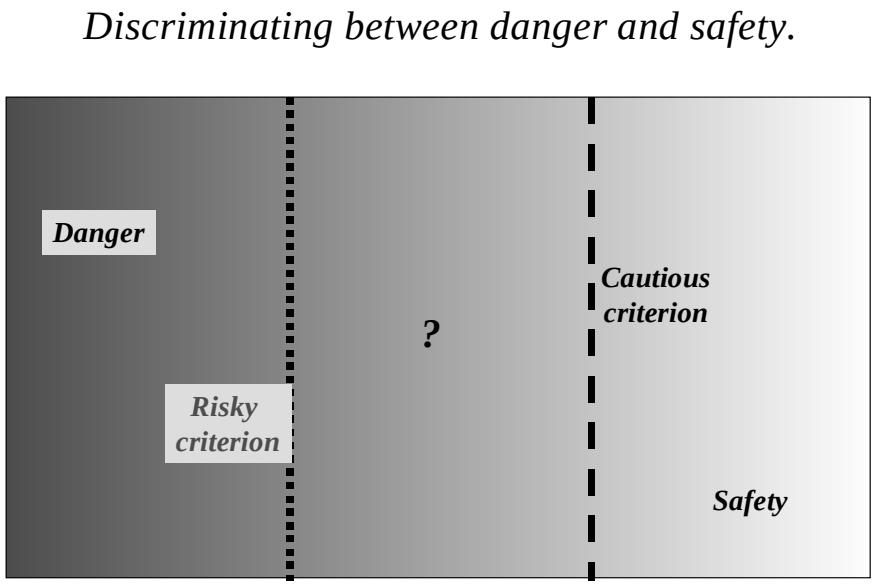

Another way of thinking about this is in terms of a diagram where one axis represents the true properties of the object (e.g. signal versus noise) and the other axis represents the perceiver's response or decision (e.g. to treat it as signal versus treating it as noise). Each of the resulting cells then has a distinct meaning. Treating a real signal as a signal constitutes a 'true positive' or 'hit'; treating a signal as noise constitutes a 'false negative' or 'miss'; treating what is actually just noise as a signal is a 'false positive' or 'false alarm' and treating noise as noise is a 'true negative' or 'correct rejection' (see Figure 2).

Figure 2. Signal detection: Decision-outcome combinations.

Action Actual risk	Treat as dangerous	Treat as safe
Dangerous	True positive (Hit) *TP*	False negative (Miss) *FN*
Safe	False positive (False alarm) *FP*	True negative (Correct rejection) *TN*
Response bias Accuracy	*[Cautious] TP+FP vs. FN+TN [Risky]* *TP+TN vs. FP+FN*	

A methodology used in SDT research is to plot the cumulative frequencies of these different responses over time for a single decision-maker, with the vertical axis representing the proportion of hits and the horizontal axis the proportion of false alarms. The result is termed the ROC ('Receiver Operating Characteristic') curve. Figure 3 shows two such curves, the upper one that of a person or system with high sensitivity, the lower one with low sensitivity. The area under the ROC curve is a measure of sensitivity. A straight diagonal line from the lower left to the top right would represent discrimination at chance level. A curve that exactly traced the left vertical and the top horizontal would represent perfect discrimination, 100 per cent correct responding. Also shown are the projections for two different response criteria. The more cautious criterion produces a lower probability of misses (shown where the horizontal dashed lines meet the right vertical), but a higher

probability of false alarms than the risky criterion. Note the effect of differences in sensitivity on the proportion of false alarms produced by either criterion (shown where the perpendiculars meet the bottom horizontal). In particular, the adoption of a cautious criterion under conditions of low sensitivity produces a very high proportion of false alarms (see rightmost perpendicular). This may be analogous to the adoption of a precautionary principle or approach under circumstances when there is a lack of evidence about whether something is truly safe or dangerous.

Figure 3. Receiver operating characteristic curves.

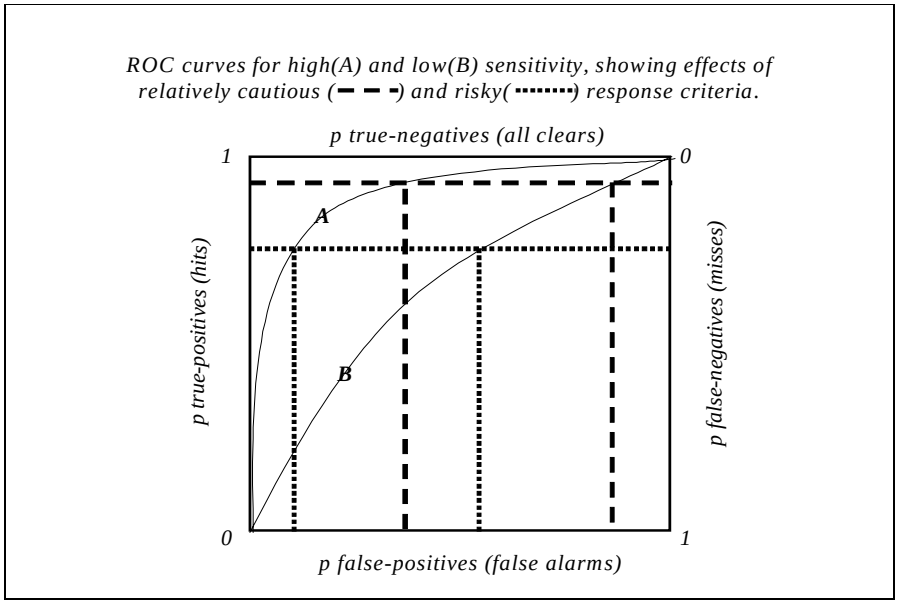

So far, this framework distinguishes between different types of errors, but says nothing about their possible costs. Where a precautionary approach is adopted, this is usually because the potential cost of a miss, perhaps in the shape of a risk to the environment or public health, is perceived to be high, and not sufficiently offset by the benefits of the activity or product continuing unchecked in the absence of better evidence of its safety. On the other hand, if tests for danger are believed to be really precise (sensitivity is high), the probability of misses can be reduced to a very low level without too many false alarms. Safety-critical industries typically aspire to design ‘defence in depth’ against the potentially catastrophic consequences of error or system failure.

Consider how this might apply in the context of medical diagnosis (Figure 4). Let us assume that the treatment of choice will indeed benefit the patient if he or she really has the disease, whereas leaving the disease untreated will have serious consequences. This means that a true positive diagnosis is much preferable to a false negative. But suppose that the consequences of a false-positive diagnosis (inappropriate treatment when no disease is present) are

even worse than simply leaving the disease untreated. If this were the case, this should make the doctor much more cautious about making a positive diagnosis and prescribing treatment.

Alternatively, treating the patient when the disease was absent might have no serious side effects. This should produce a bias towards treating them rather than not.

Where the benefits and costs of different correct and incorrect decisions are known in advance, we can set response criteria that will be more or less 'rational', depending on how accurate this knowledge is and how equitably the interests of different parties are considered. But this is only half the story as far as risk perception and risk-taking are concerned. The benefits and costs accruing from different decisions are themselves feedback that will modify the decision-maker's expectancies and future behaviour, and not always in the direction of greater accuracy. To illustrate some of the learning effects that might result from the doctor's different decisions in the above example, I have added a further row, labelled 'learning,' to the figure.

Suppose, for example, that over-treatment, or unnecessary prescription, had no adverse effects, and that after a while the wrongly diagnosed patient's symptoms disappeared anyway. This could lead to illusory over-confidence in the treatment's effectiveness. "When I prescribe this treatment, the patient always gets better". One of the difficulties with medical decisions based on personal clinical experience (as opposed to controlled clinical trials) is that individual doctors may not receive sufficient feedback about what happens in all four cells of the matrix. It is difficult for the individual doctor, on the basis of a few cases, to draw reliable counterfactual inferences – how one patient might have fared even if untreated, or how another might have fared if he or she had received a treatment that had been withheld or was unavailable at the time.

Figure 4. Decision-outcome combinations for medical diagnosis.

Diagnosis Actual disease status	Disease present	Disease absent
Disease present	True positive Correct diagnosis *Appropriate treatment*	False negative Incorrect diagnosis *Failure to treat*
Disease absent	False positive Incorrect diagnosis *Inappropriate treatment*	True negative Correct diagnosis *No unnecessary treatment*
Learning	Not necessarily *Unless inappropriate treatment harms patient, improved health increases confidence in treatment.*	Hopefully *Poor outcome following failure to treat should prompt test reappraisal.*

Tendencies such as false alarm expectancies and risk aversion may be a consequence of the incomplete feedback provided by everyday experience. The approach-avoidance decisions considered earlier are illustrated in Figure 5. (For consistency with the previous examples, actions are labelled ‘positive’ when they involve avoidance, i.e. judgement that danger is *present*, even though the associated attitude or evaluative expectancy will be negative.) As this shows, other things being equal, a decision-maker has no way of knowing that a decision to avoid was correct or incorrect, since no feedback results. Hence, whereas incorrect approach may be corrected by experience, incorrect avoidance will not.

Figure 5. Decision-outcome combinations for approach and avoidance.

Action Valence	Avoid	Approach
Bad	True-positive *Correct avoidance, but no feedback*	False-negative *Reduced approach due to punishment*
Good	False-positive *Prejudices, phobias persist in absence of feedback*	True-negative *Strengthened approach due to reward*
Learning	No *But absence of harm may reassure.*	Yes *Feedback improves discrimination*

Trust in the future

In many important situations, feedback is not immediate, and expectancies concerning short-term and long-term consequences need to be balanced against each other. Psychological research on learning in animals demonstrates a very general principle: immediate rewards have more impact than delayed rewards of the same magnitude. This means that if rewards (or punishments) are delayed, they have to be bigger to have the same effect. Adapting the proverb, a bird in the hand is not just more certain than two in the bush, it is here and immediate, and does not need waiting for.

In economic terms, animals appear to use a form of 'temporal discounting' whereby they attach a lesser utility to outcomes that may occur some time in the more distant future. Similar processes occur with humans. Several studies point to the difficulties that young children in particular have with 'delay of gratification'. Faced with a choice between a small reward now, and a large reward after some delay, many children find it difficult to wait (Mischel, 1974; Mischel & Mischel, 1983). Nor are adults immune. Several health and lifestyle problems are associated with behaviour such as smoking or extravagant expenditure which produces immediate short-term pleasure at the cost of some bad effect, such as disease or debt, in the longer term.

A number of factors may be associated with the tendency to base decisions on short-term outcomes. One may be the ambiguity of the chain of cause and effect to the ordinary perceiver, especially of actions whose effects may become apparent many years later. Several health outcomes provide examples. They arise in different forms depending on whether the cause is more likely to be cumulative (e.g. smoking and lung cancer, or poor diabetic control and retinopathy) or due to a specific accident or infection (e.g. mesothelioma from asbestos exposure, cancers linked to radiation accidents, malaria, sexually transmitted diseases). Without an adequate mental model of the relevant disease process, individuals may remain unaware of any link between their health outcomes and their own behaviour.

As well as ambiguity looking backwards, there is unpredictability looking forward. A promise of 'jam tomorrow' may be all very well, but tomorrow, or at least the jam, may never come. Why give up the 'sure thing' of an immediate reward for the possibility of a better outcome far in the future? This is our familiar principle of risk aversion in a new guise. Many behaviours (especially self-indulgent habits that carry a long-term health risk) that are defined as 'risky' on medical grounds may be risk-averse, using the term in its strict technical sense, in terms of the individual's perspective. Delay of gratification starts becoming worthwhile if one regards the future and its associated rewards as predictable. The capacity to forego immediate temptations for a longer term goal is not just something we gain through maturity. It is also something we learn through finding out that we can, within limits, control the outcomes we receive as a consequence of our own behaviour.

Social learning theory is the name given to the field of psychology concerned with how we develop distinctive cognitive and behavioural styles based on our

interactions with one another and the good and bad things that happen to us as consequences of our actions (Mischel, 1973; Mischel & Shoda, 1995). A variety of terms describe different aspects of this process and the personality differences to which it is claimed to give rise. Rotter (1966) proposed that individuals differ in terms of their 'locus of control'. Those with 'internal locus of control' believe that they have more personal control over events, whereas those with 'external locus of control' believe that their fate is in the hands of other people or chance. Bandura (1977) proposed the similar concept of 'self-efficacy' to account for individuals' general confidence in their own ability to engage in difficult tasks (e.g. quitting smoking). At the other extreme, Seeman (1971) related socio-political alienation to learnt feelings of powerlessness, while Seligman (1975) introduced the notion of 'learned helplessness' to account for aspects of clinical depression.

The unifying theme in all these approaches is that, as a result of the good and bad things that happen to them, people acquire expectancies about the world that guide their behaviour and affect their views of themselves. At one extreme, they may come to view the world – or more accurately the 'home range' that they personally inhabit – as a predictable and mostly friendly environment. Here they can invest their time and energy for a better future in the confidence that they have reasonable control over their own destiny. They can bring their plans to fruition and receive the rewards, social approval and feelings of self-worth than they deserve. At the other extreme, people may view their environment as unpredictable, hostile, and beyond their control, threatening them with dangers and degradations that they are powerless to prevent.

Sources of evidence

The focus in most of the work I have so far reviewed has been on the individual decision-maker, forming expectancies and making choices on the basis of evidence and feedback. Less has been said about where this evidence comes from. The simple case is where we gain evidence from direct experience. In other words, things happen to us personally that are good or bad, and they confirm or contradict our expectancies. But, as social creatures, we are not restricted to evidence from our own experience alone. We can observe what happens to other people. We can draw inferences from the consequences of their actions about what would happen to us if we behaved similarly, and we can feel pity or envy at their plight or their success. This can be termed vicarious experience. We can gain such experience not only from actually being there, but also through television, with its vivid and often live portrayals of events. Such experience is less direct than first-hand observation and the portrayals reflect decisions made by programme editors, but the broadcast images, for example of the destruction of the Twin Towers, the war in Iraq or victims of famine or AIDS in Africa, can still have an immediate and emotional impact that most people could not experience otherwise.

Observing others also provides the opportunity for imitation. Since imitative behaviour is widespread in other species and readily observable in human

infants, it doubtless has instinctual underpinnings and has proved functional during evolution. At the level of more deliberate decision-making, it can take the form of following the example of those we see as admirable or successful.

Even if our own experiences and observations provide an insufficient basis for decision-making, however, we can turn to other people who are (or claim to be) better informed than we are. The use of indirect experience – what other people know and can tell us – is one of the most important aspects of risk in general and public perception of risk in particular. Our ability to communicate with each other has contributed to our survival throughout evolution. Among many other advantages, it has enabled our ancestors to inform each other about sources of food, to warn each other about dangers, and to coordinate cooperative activities in pursuit of shared goals. It also has allowed the development of cultures and societies based on principles of distributed knowledge, whereby different members can have different forms of expertise (rather than everyone individually having to know everything) but can share this expertise with each other when called upon to do so. Much the same principles underlie contemporary notions of distributed processing within cognitive systems.

The relevance of this to risk perception is that, whenever we are considering innovations or changes beyond what we ourselves have experienced or observed, it is natural for us to ask other people for information, advice and perhaps protection from possible harmful consequences. We seek to supplement our own direct experience with indirect experience, and in so doing make ourselves dependent on others. This dependence can take the form of relying on others as sources of information about events that are beyond our own experience, but often the same ‘experts’ who provide such information also take decisions on our behalf about how best to manage risks that are beyond our personal control. These forms of reliance both require trust, in broadly similar ways but with somewhat different consequences for behaviour.

Trust in others as decision-makers

There are many examples of our reliance on others to manage risks on our behalf. We put ourselves in the hands of doctors for medical treatment. When we get on an aeroplane, we are trusting a whole system of professionals, including pilots, manufacturers, maintenance staff, and air traffic controllers. Deciding that air travel is safe involves an implicit or explicit judgement of trust in all of them. Just take air traffic controllers, for the sake of simplicity. We need to trust the air traffic controllers to make good decisions. In other words, we need to feel that they will be able to tell the difference between safe and dangerous situations, that they will not let us take off unless they judge it to be safe, but, by the same token, will not keep us on the ground unnecessarily. In other words, we expect air traffic controllers to exhibit both good discrimination ability and an appropriate response criterion.

Signal detection theory distinguishes the parameters of discrimination ability or sensitivity, and response criterion or bias, to provide a basis for evaluating

a decision-maker's performance. When we trust or distrust some person or agency to make decisions and to control risks on our behalf, we are effectively evaluating that person's or agency's performance. This makes it easier to see how trust can depend on, or be undermined by, different kinds of evidence. It would be no comfort to know that an air traffic controller, doctor or power station manager had our safety at heart if we did not feel they were also sober and vigilant and had been adequately trained to distinguish safety from danger. Likewise, it would be worrying to feel that any of these professionals, however competent, would let their decisions be inappropriately biased by financial concerns (e.g. to keep gaps between flights below a safe minimum, to prescribe medications on the basis of manufacturer's incentives, to compromise on safety inspections).

But all this raises a by now familiar difficulty. How can I personally find evidence of another decision-maker's discrimination ability or bias? The classic signal detection paradigm requires calculations based on the outcomes of a whole series of decisions, and with full feedback to the person making the calculation about whether these decisions were correct or not. But if I were to ask "How good is my doctor?" I have nowhere near this kind of direct evidence. I have experience of only a very few consultations. My health may have improved after such consultations, but it might have done so anyway, without any help from my doctor. If I avoid the bad experience of a clear misdiagnosis, say, I am likely to take my doctor's ability more or less for granted. In much the same way, I will take the competence of pilots, air traffic controllers, taxi drivers and school teachers for granted unless I have a specific reason not to. However, I typically have neither the resources nor the inclination to monitor how competent these professionals are over a series of decisions, let alone to make any comparative judgements of how good their decisions by the normal standards of their profession. If anything bad happens, I may notice, but otherwise a normal consultation, flight, taxi ride or school day is effectively a non-event that requires little or no change in my view of the world.

This means that misses with bad consequences attract our attention (and frequently the attention of news reporters), they become news, and prompt us to reappraise our view of the people or systems that have produced them (Rozin & Royzman, 2001). Earlier we saw how feedback concerning the consequences of different choices could lead to learning on the part of a decision-maker, and how incomplete feedback could lead to different kinds of bias, including risk aversion. Something very similar can happen with perceptions, since a perceiver (e.g. the public) will also be learning about the performance of the decision-maker or system. And again, such learning can lead to biased perceptions as a consequence of incomplete or ambiguous feedback.

Feedback can be especially incomplete and ambiguous when we are considering systems or circumstances that are, in purely statistical terms, very safe. This does mean that such systems contain no danger and involve little need for vigilance. Take driving for example. Most car drivers, apart possibly from the pathologically anxious or those who have recently had an accident,

start their journeys with every confidence that they will reach their destination unscathed. Up to a point, such confidence is well grounded. By far the great majority of car journeys involve no accident. As a consequence, most of us think we are good drivers. But accidents most certainly do happen, even to drivers who think they are pretty good. It's just that hazards that are apparent with larger samples of data, accumulated over long periods of time and many people, may be undetectable at the level of the individual case.

Closely related to this is the fact that many systems and activities can (again in purely statistical terms) be relatively 'forgiving' or tolerant of unsafe practices. We don't only avoid accidents when we drive well. Usually we remain safe even when we drive badly. Drivers may speed, ignore traffic lights and even consume moderate amounts of alcohol and still arrive safely. Drivers who do these things are taking a far higher risk (for themselves and others) than those who don't, but with luck and favourable road conditions, they may, perhaps more often than not, escape unscathed on any given journey. Likewise cigarette smoking is anything but safe. There are many different health consequences to consider, but just regarding lung cancer, smokers are hugely more at risk than non-smokers. Yet still, 'only' about 1 in 10 smokers will get lung cancer. Take this into the industrial context. Poor maintenance and operational practices increase the risk of accidents, whether on railways, oil rigs, nuclear power stations, surgical wards or fishing vessels.

With a large enough recorded sample of incidents, an increased statistical risk can translate itself into someone getting hurt (or worse). But if accidents are rare, by definition they will not occur in most small samples of recorded performance. Small data sets can be misleading, and unrepresentative of the population from which they are drawn. Kahneman & Tversky (1971) use the phrase 'the law of small numbers' to refer to people's readiness to over-generalise from small sets of data. If we have only a small sample of instances from which to judge the safety of a system, and this sample does not include an accident or adverse event, we are likely to falsely conclude that the system is safe. Much of the time, feedback does not unambiguously distinguish 'correct rejections' (true negative beliefs that a safe system is safe) from 'misses' (false negative beliefs that a dangerous system is safe). Potentially catastrophic flaws may remain undetected or be discounted by the system's operators, while the public remain in blissful ignorance of the dangers to which they are subject.

The other side of the story is that if an accident does happen, we can over-generalise in the other direction. Sometimes an accident can happen even if all reasonable and affordable steps had been taken to keep risk to a minimum. Even good drivers can have accidents. With hindsight, different precautions might have been taken, but not necessarily better precautions. Risk A might have been knowingly accepted because, in the long run, this meant that the more serious Risks B and C could be avoided. But this is little comfort if Risk A actually happens, and we are on the receiving end. A treatment may have unwelcome side-effects for some patients, but this does not make it the wrong treatment, even for those patients (though obviously it would even better to be able to identify in advance which patients might

experience most side-effects and to attempt to devise a different treatment for them). A basically safe procedure, or a whole industry, may be rejected on the basis of an isolated negative event.

But in many cases, even an 'isolated' disaster, for instance a nuclear meltdown, is one disaster too many, and the public has every right to assume that none will occur. But does Chernobyl (even) show that nuclear power is 'basically' unsafe? This kind of question cannot really be answered without a more detailed analysis of how that accident occurred, and whether the determining factors were common to all forms of nuclear technology or specific to particular reactor designs and modes of operation. The problem faced by advocates of nuclear power is that such technical distinctions can seem like splitting hairs to members of the public whose attitudes to nuclear power are less differentiated. To take a possibly less controversial example, should news of an accident involving a cut-price charter air-line from a developing country make one more nervous of flying with, say, British Airways? 'Over-generalisation' in such contexts consists of adopting a single attitude towards too wide a category of industrial practices and circumstances. Such over-generalisation can be part of what is sometimes termed 'stigmatisation'. As defined by Kasperson, Jhaveri and Kasperson (2001, p. 15), this is "the process by which individuals select an attribute of a person, place, technology, or product and denigrate the possessors of that attribute, discriminate against the possessor, and may even construct a stigma 'theory' or 'story' to explain the inferiority and its roots."

Generalisation comes about whenever perceivers evaluate complex objects on the basis of only part of the full story. Since we rarely have the cognitive capacity to give due weight to all sides of a question, most evaluations involve selective processing of information. It can take the form of conscious inductive reasoning or prediction (as when a punter studies the recent form of runners in a horse race), but it can also be driven by learnt associations that have become so automatic (Bargh, 1997; Fazio, 2001) that we may be unaware of their origin and their impact on our judgements and preferences. As pointed out by Slovic et al. (2002), affective and emotional associations have a strong potential to influence risk judgements in this way.

Trust in others as communicators

Reliance on others as sources of information, or as communicators, has been extensively researched, though not necessarily from the standpoint of risk. The conclusion is that, under certain conditions, our own opinions can be strongly influenced by what other people think, or what we think they think. This can even happen with simple perceptual judgements, in experiments where confederates of the researcher deliberately give wrong answers before the participant responds (Asch, 1956). A sizeable minority of participants copy the confederates' obviously incorrect responses and doubt the evidence of their own eyes.

Although it is often thought to suggest weakness or irrationality, conformity to others' points of view can be a reasonable and functional strategy for interpreting events in an uncertain world. If something is 'real' and 'objectively true', there should be agreement over the facts of the matter. Under conditions of uncertainty it often makes good sense to look for a consensus of opinion as a guide to what is actually the case. The Asch (1956) experiments work because they violate this assumption, or more precisely exploit it, by confronting participants with an apparent consensus opposed to their own perception. When this apparent consensus is undermined by having one of the confederates be an ally to the participant and provide a correct answer, the participant is much less likely to conform to the majority opinion.

But a strategy that is adaptive under many conditions can produce big problems in others. One is the situation when we seek the opinions of others who are no better informed than we are ourselves. We may discover a consensus, but it may be utter myth rather than a guide to the truth. Yet the very fact that others believe the myth will make it more likely that we will too. Indeed, even without any such encouragement from other people, we may overestimate the extent to which our own opinions are commonly shared. At least as often, however, different people may express different opinions. What do we do? We have to make choices about whom to believe and trust, and whose opinions to discount or ignore. One of the simplest and most powerful principles involved is that we are generally more prepared to accept the opinions of other people whom we like than those whom we dislike. This is enough of a potential bias in itself, but it combines with the following, equally powerful, principle: we tend to like others more if they share our own opinions than if they disagree with us (Heider, 1946; Newcomb, 1953). Together, these principles set up a positive feedback loop whereby we seek confirmation of our beliefs from the very people least likely to challenge what we think, and also come to value their opinions even more because they provide the support we seek. At the same time they may be seeking the same kind of confirmation from us (Eiser, Claessen & Loose, 1998). An extreme example of this, termed 'groupthink', refers to the tendency of authoritarian and over-cohesive groups (such as those surrounding some political leaders) to disallow minority opinions and resist discussion of evidence that threatens their established point of view or strategy (Janis, 1972; Raven, 1974).

So we rely a great deal on other people's opinions to help resolve uncertainty. The extent of this reliance will depend on our personal relationships with them, but also on other aspects besides. Not infrequently, information comes to us from people we don't know personally, but who claim to be knowledgeable or expert in some way. Public perception of risk often comes down to a matter of how the public views judgements about the seriousness of a hazard communicated by government, scientists, the media and other agencies. These judgements can be seen as yet another example of a signal detection decision, and the same questions can be asked about the accuracy and bias of this decision. How much reliance should be based on a newspaper article critical of GM crops, for instance? Is the journalist who wrote it credible? Adopting a signal detection perspective, this question splits into at least two sub-questions. First, does the journalist have any special knowledge about the issue? In other words, would he or she be able to tell the difference between safer and more dangerous crops? Second, does the journalist have a bias towards interpreting any evidence as favourable or unfavourable to GM? In other words, is the journalist sufficiently expert and unbiased? If not, little reliance can be put on the information or arguments presented in the article.

But in the context of communication, we need to ask a further question. Where does any bias lie? Is it at the stage of the interpretation of the information, as when a doctor looks at the results of a diagnostic test, or at the stage of communication, as when a doctor decides how to explain the results to the patient? With doctor-patient communication, we now have a cultural norm of openness and honesty. Doctors are expected to give their patients full information, even if the news is bad. But giving bad news is difficult, and it is not that long ago that many doctors felt justified in withholding information that they felt would cause their patients distress. Such doctors could be said to have had a communicative response bias in the direction of playing down the seriousness of the patient's condition.

With trust, we are looking at perceptions, and just as we might seek cues about the competence and regard for safety of an airline a or rail maintenance contractor, we can ask what kinds of events or circumstances might cast doubt on a communicator's expertise and honesty. Of special importance could be evidence concerning what, if anything, the communicator stands to gain or lose by dressing up the information in a particular way. Expertise by itself is not enough to engender trust. Independence and disinterestedness are also essential. If biomedical scientists employed by the tobacco industry try to tell us that smoking is safe, or non-addictive, our reaction is likely to be "Well, they would say that, wouldn't they?" Wherever we rely on 'experts' to tell us whether something is safe or dangerous, we need to know not only "Can they tell if it will have bad effects?" but also "And would they tell us anyway?"

Any communicator who is perceived, rightly or wrongly, to have a stake in persuading us either that something is safe (for instance if they are paid by the manufacturer) or that it is dangerous (for instance to sell newspapers) is less likely to be trusted than someone who is seen to provide the facts as they

themselves interpret them. This can be represented in terms of Figure 6. The four categories of decisions distinguished by the signal detection framework lend themselves to different epithets. True-positive communications about the presence of a risk are ‘vigilant’, false-negatives are ‘complacent’, false-positives are ‘alarmist’ and true-negatives, ‘reassuring’.

Figure 6. Decision-outcome combinations for risk communication.

Communication Actual risk	Declare to be dangerous	Declare to be safe
Dangerous	True-positive *Vigilant*	False-negative *Complacent*
Safe	False-positive *Alarmist*	True-negative *Reassuring*
Communication bias (readiness to declare presence of risk)	*[Quick] TP+FP vs. FN+TN [Slow]*	

Social amplification

This points to the importance of the formal and informal social networks through which information and influence flow. A highly influential descriptive framework with relevance to risk is Kasperson's model of social amplification and attenuation (Kasperson & Kasperson, 1996; Pidgeon, Kasperson & Slovic, 2003). The basic notion is that ideas about risks are transformed as information is transmitted through different sources and channels of communication, on the analogy of a broadcast radio message. This transmission process, or 'information' flow, starts with 'sources of information' (personal experience, direct and indirect communication) concerning risk and risk events. This information then passes through 'information channels' (e.g. informal social networks) to 'social stations' (e.g. media, government agencies or action groups). Individuals then bring their own mental processes to bear on the information ('individual stations') which then guide changes in institutional and social behaviour, perhaps by causing attitude change or

protest. Expressions of opinion and actions by individuals can then have 'ripple effects' outwards from more immediate to more distant networks of social relationships (e.g. the local community, professional groups, and through to wider society) with a variety of economic, political and environmental 'impacts' (e.g. financial loss to industry or blight on property values).

Where all this leads to heightened concern, this is termed 'social amplification', where it leads to reduced concern, 'attenuation'. The process of information flow can involve feedback and repetition (individuals may seek the views of neighbours and family as a result of how they interpret a government or industry statement). Implicit in this notion of feedback is the idea that individuals may sample information selectively with a bias towards *confirmation* of their initial impressions, since like-minded sources of information are more likely to be listened to and trusted. Combining what Kasperson suggests about information flow with what we know from elsewhere about information search implies that there is a self-reinforcing dynamic implicit in the whole process of information transmission. The likely effects of this are a reduction of uncertainty for the individual and a polarisation of opinion into contrasting camps within society as a whole.

The stigmatisation of particular industries or neighbourhoods can be a special case of this, where individual antipathies become socially reinforced and adopted as cultural truisms.

Dynamics of change and resistance

Risk perception is a response to uncertainty. But one of the most striking things about people's opinions about risks is that they often appear to be held strongly and acted upon with apparent certainty. We need this certainty to exert control over the good and bad outcomes we receive and this leads us to push at the boundaries of what we know. But seeking information can be costly in terms of time and effort. It can also be dangerous. Animals which explore their environment in search of new sources of food can put themselves in real physical danger, no less than the European explorers who set sail five hundred years ago in search of new continents beyond their horizon. The rewards may or may not outweigh the costs. Seeking information also carries psychological danger. We may well find out something we didn't really want to know, or spend more time and effort reconsidering – or regretting – a decision that we thought was done and dusted. Under more extreme circumstances, we might be forced to surrender cherished values and allegiances. So long as things are basically alright as they are, why not let sleeping dogs lie?

Despite this inhibition against unconstrained exploration, learning still occurs and leads us to adjust our evaluations of things. The big trouble is that such learning can be highly selective. Other things being equal, we are less likely to engage with things we initially anticipate to be dangerous than with things we anticipate to be safe. Hence, false-alarm reactions, prejudices and unfounded fears may persist for want of the new information and direct experiences that

could convince us that we were wrong. At an individual level, our attitudes guide our behaviour, and our behaviour largely reinforces our prior attitudes. At a social level, we let ourselves be guided by those we trust, and we are less likely to trust those whose viewpoints or motivations seem at variance with our own. So our networks of sources for advice, information and influence also become strengthened through repeated use. Attitudes and risk perceptions can be self-reinforcing and self-fulfilling and once formed, can be held in place by large networks of associations, both in the mind of the individual and via social relationships of affection, identification and trust.

But change can still occur. What kind of change can we expect within a system characterised by the self-reinforcement of divergent positions? One of the common features of dynamic systems is that, depending on the starting point, the same amount of information can produce either a large or small change, or sometimes no change at all (Stewart, 1989). Think of the relationship between rainfall and flooding. Up to a certain level (the capacity of a flood defence scheme, or natural river bank), rainfall can increase with no appreciable risk of flooding. Beyond that level, a single storm can lead to a river bursting its banks and large areas being inundated. In terms of human attitudes, it has been observed that it is far more difficult to persuade somebody to 'change sides' from moderate support to moderate opposition towards some government policy, say, than to achieve the same amount of change from moderate to extreme support (Lange & Fishbein, 1983).

Catastrophe Theory (Zeeman, 1976) is one of a family of models of nonlinear dynamics and offers one of the better known descriptions of how sudden change can occur under some conditions but not all. The basic idea is represented graphically in Figure 7. Imagine we are predicting changes in an output – say, belief about some issue, perception of risk, or trust in a communicator – on the basis of two inputs. Belief change can be mapped as a trajectory over a plane defined by the two parameters, but in the special case illustrated, this plane is curved and partly folded. What this means is that, depending on the course of the trajectory, change will be more or less difficult and may even involve a sudden jump. The figure shows possible relationships between a change in belief (risk perception or trust) and two control variables: the amount of information provided, and the individual's involvement in the issue (see van der Maas, Kolstein & van der Pligt, 2003 for a fuller technical account of this example). 'Involvement' refers to the perceived importance of the issue and overlaps with the concept of 'attitude strength' (Petty & Krosnick, 1995). It will tend to be high if the issue seems to be strongly associated with other important values, or closely tied in with a person's social relationships and identity as a member of a given group. Such associations will themselves reflect prior learning and will tend to inhibit certain kinds of change in belief.

Figure 7. Catastrophe theory illustration of possible impact of information and involvement on belief and risk perception.

Involvement

Under conditions of relatively low involvement (the back of the plane) an increased amount of information (such as health warnings) leads to a one-way change in belief (perhaps a higher perception of risk) from point D to point E. However, under conditions of high involvement, an equivalent change in belief from A to C would not only be more difficult (the individual could get stuck at B), but if it occurs, could be marked by a sudden jump or step change such as deciding to quit smoking. Conversely, a change from C to A could be easier and more or less irreversible, as when trust is lost in a previously respected expert. Gradual change, sudden change and resistance to change can all be observed within a single system, depending on the conditions.

Implications for decision-making

Recognising that we are prone to bias and errors of reasoning is one thing, but knowing what to do about it is quite another. Recognising our own fallibility is a start, but only a start. The problem is that we don't typically know when we're making a mistake until afterwards, and sometimes not even then. Even bad decisions can feel right. This is partly because we acquire our beliefs, values and feelings through experience. This process is highly adaptive, and essential to our survival. The trouble is, it is also highly selective. Relevant facts can be ignored, and particular outcomes or contingencies can be under-sampled. Furthermore, it is extremely difficult to unlearn habits of thought and action that have been built up over a long time, or even introspect critically about why we feel the way we do, when such feelings appear so automatic. Our previous experience has 'got us here'; on the other hand, we generally have very little insight into how we got here or what we have missed out on in the process.

Our capacities for thinking and feeling have evolved in the way they have because of their benefits for survival. This isn't to say that we never have dysfunctional thoughts or feelings, nor that all the instincts that helped our ancestors are necessarily adaptive in the modern world. However, our feeling and thinking serves a purpose, that of helping us achieve desired goals. Being able to predict events improves the possibility of controlling what happens, and hence the probability of attaining our goals. It helps if our predictions are accurate, but accuracy is not an end in itself. Our decision-making capacities are not geared up to provide us with 'the truth, the whole truth, and nothing but the truth' for its own sake, but to help us achieve desired goals and avoid undesired ones. Research on risk perception has tended to be predominantly concerned with judgements of probability. But probabilities only turn into risks when they relate to things we want to avoid, achieve or obtain. Wants can be reasonable or unreasonable and ethical or unethical, but they can't be accurate or inaccurate.

From the perspective of decision-making, however, the danger is that our wants can lead us into forms of inaccuracy of which we're unaware. We don't know what we don't know, and find it difficult to imagine how things could be otherwise than as they appear to us. Often too there are many things we don't want to know, particularly if they might make us upset or confront us with distressing or difficult decisions (Janis & Mann, 1977). So how can we

safeguard ourselves against this danger? Ideally by being less selective in our information-processing, but this is easier said than done. More practicably, by being open to advice from others, since different people may look at different information in different ways, and have experienced cases and situations that we have not explored. But openness to new advice is a fragile quality, vulnerable to the countervailing forces of conformity and commitment to prior beliefs. Organisations, no less than individuals, tend to be risk-averse if they fail to allow space for at least some individuals to engage in innovatory exploration of potentially costly but also potentially beneficial alternatives. Innovation is likely to be facilitated in organisational cultures that tolerate diversity and avoid premature closure on dominant viewpoints (Denrell, 2003; Denrell & March, 2001).

Implications for risk communication

Research on risk communication, like that on risk perception, has largely been concerned with how best to convey probabilities, or more simply, 'facts' about the presence of absence of some hazard or vulnerability. Assessments of the success or failure of risk communication have likewise tended to focus on the degree of agreement between the warnings or reassurances conveyed by communicators and the degree of concern expressed by their audience. Health education is deemed to be relatively successful if it leads to young people who have a poor diet becoming more aware that they may be increasing their risk of obesity, heart disease or cancer. Sometimes there is an attempt to communicate quantitative statistical risks, where these can be estimated with reasonable confidence, but often it is simply a matter of trying to inform the audience that a risk does or doesn't exist. The subsequent 'accuracy' of the audience's risk perceptions can usually be measured only very crudely, if at all.

One of the advantages of this approach is that the communicator can claim to be a messenger rather than a propagandist, stating the 'facts' and letting the audience draw its own conclusions. Often this claim is sincere, but sometimes it is disingenuous. Either way, it needs to be viewed with scepticism. Some element of selectivity is almost bound to enter in, both on the part of the communicator (deciding what it is that the audience needs, or should be allowed, to know) and on the part of the audience (deciding which bits are most important and/or understandable). It is probably a vain hope to think one can always provide complete information in a manner that will be completely and immediately understood. Even when complete disclosure, perhaps of the risk of possible medical side-effects, is obligatory, additional advice may often need to be given about which risks are the most important. However the alternative, deliberate incomplete disclosure, is even more dangerous. If discovered, it can destroy trust while complete disclosure, especially of evidence potentially damaging to the communicator's own interests, facilitates trust.

Selectivity of how information is interpreted by the audience is also a fact of life, and can be unwittingly or manipulatively exploited by communicators by virtue of how its possible consequences are described, as research on

'framing' shows. Events described in more detail can be seen as more likely (Tversky & Koehler, 1994). Rothman & Salovey (1997; see also Detweiler et al., 1999; Devos-Combey & Salovey, 2002; Rothman (Rothman et al., 1999) studied the effects of framing on the effectiveness of health messages, arguing that messages framed in terms of gains (ways to improve health) tend to be more readily accepted than those framed in terms of losses (warnings of damage to health).

Risk inevitably involves a concern with good and bad outcomes. Such values will be implicit in any exchange. What any audience will be listening for is an indication of what they need to do – not simply for a prediction of what consequences may or may not happen, but also for an indication of how good or bad these consequences will be, and what they, or anyone else, is expected to do about it.

So how are communicators to avoid falling into the role of propagandists? A large part of the answer is to adopt a less didactic model of communication. It is already widely recognised that, to be effective, the communication of risks cannot be merely one-way. It must involve exchange and interaction between all parties. But there are obstacles to putting such good intentions into practice. The 'experts' need to be prepared to give up not only time, but some of their power. The 'public' need to be willing to be engaged in the decision process and take some responsibility for the outcomes, if these have been shaped to take account of their views. All this can be costly, in terms of time, patience and the risk of disappointment if not everything one wants can be achieved. Yet the real advantage over the traditional, less consultative, approach is that the debate or discussion can focus on what people want to know. And this can involve considering issues of value rather than merely probability.

One-way communication that fails to address issues of concern to the audience will fail to persuade, reassure or even inform properly. An example is provided by the so-called 'consultations' undertaken by the UK nuclear industry in the early 1980s to reconcile local residents to the idea of new nuclear power stations in their neighbourhood. Considerable effort was put into reassuring people that the risk of any escape of radiation was remote. But more tangible issues of concern to residents, such as damage to the landscape and disruption to the community, were inadequately addressed (Barnes, 1990; van der Pligt, Eiser & Spears, 1986; Eiser, van der Pligt & Spears, 1995).

Risk perception and experience: a summary model

The main processes whereby risk perceptions come to reflect experience and feedback are summarised in the following flow-chart.

It is assumed that risk perception starts with some event, either reported to or experienced by the individual. This event will trigger various associated memories and emotional reactions, based on the individual's previous experience of similar events. These associated memories will be collated and

will guide an initial interpretation of the event in a relatively automatic way, much as we recognise familiar objects and people. At the same time, assuming the event is public and experienced by others, it will trigger reactions from other inidividuals too, setting up the preconditions for social amplification, interpersonal influence and the recruitment of cultural values.

Depending on the amount of thought the individual is prepared to give to the issue, the initial interpretation may take on a more explicit focus concerned with estimating the likelihood and consequences of future outcomes. At this point, judgement may become influenced by cognitive heuristics and other biases, as well as by individual differences in perceived control or self-efficacy, and time perspective or the preparedness to delay gratification. At the same time, the individual may be receiving advice or information in the form of 'risk messages' from other sources, experts and so on. The impact of these messages will be mediated by how much these sources are trusted, in other words are liked, and seen as competent, unbiased and honest. These estimations and messages then feed into a summary decision process regarding whether to treat the danger as real or not.

The decision then taken has profound implications for how perceptions or risk are updated and modified by evidence. Individuals who decide to treat the danger as real and avoid exposure to it will, by and large, avoid direct harm (although they may incur costs in taking evasive action). Such an outcome is uninformative and provides no evidence either way concerning the actual presence of danger, or what would have happened if the potential risk has not been avoided. So the negative opinion that prompted the decision to avoid the risk will not be contradicted by any evidence, and may even be strengthened. This could come about if the absence of harm – which may well be accompanied by emotions of relief and such like – is over-interpreted as evidence of the correctness of the original decision.

By contrast, a decision to treat the situation as safe and approach the potential risk will provide the individual with feedback about the correctness of his or her decision. This will lead to updating of the individual's own memory of their direct experience, and modify the judgement process leading to the original decision. To the extent that that decision was based, in whole or in part, on advice from others, this can also lead to a reappraisal of the trustworthiness of these sources. Note, however, that such feedback can still be misinterpreted if it is inconsistent or delayed. Continued approach or risk-acceptance – even if imprudent – that is not followed by obvious harm is likely to be reinforced and associated with increased confidence on the part of the individual that no risk is present.

Figure 8. Flow-chart representing processing involved in judging the presence or absence of risk.

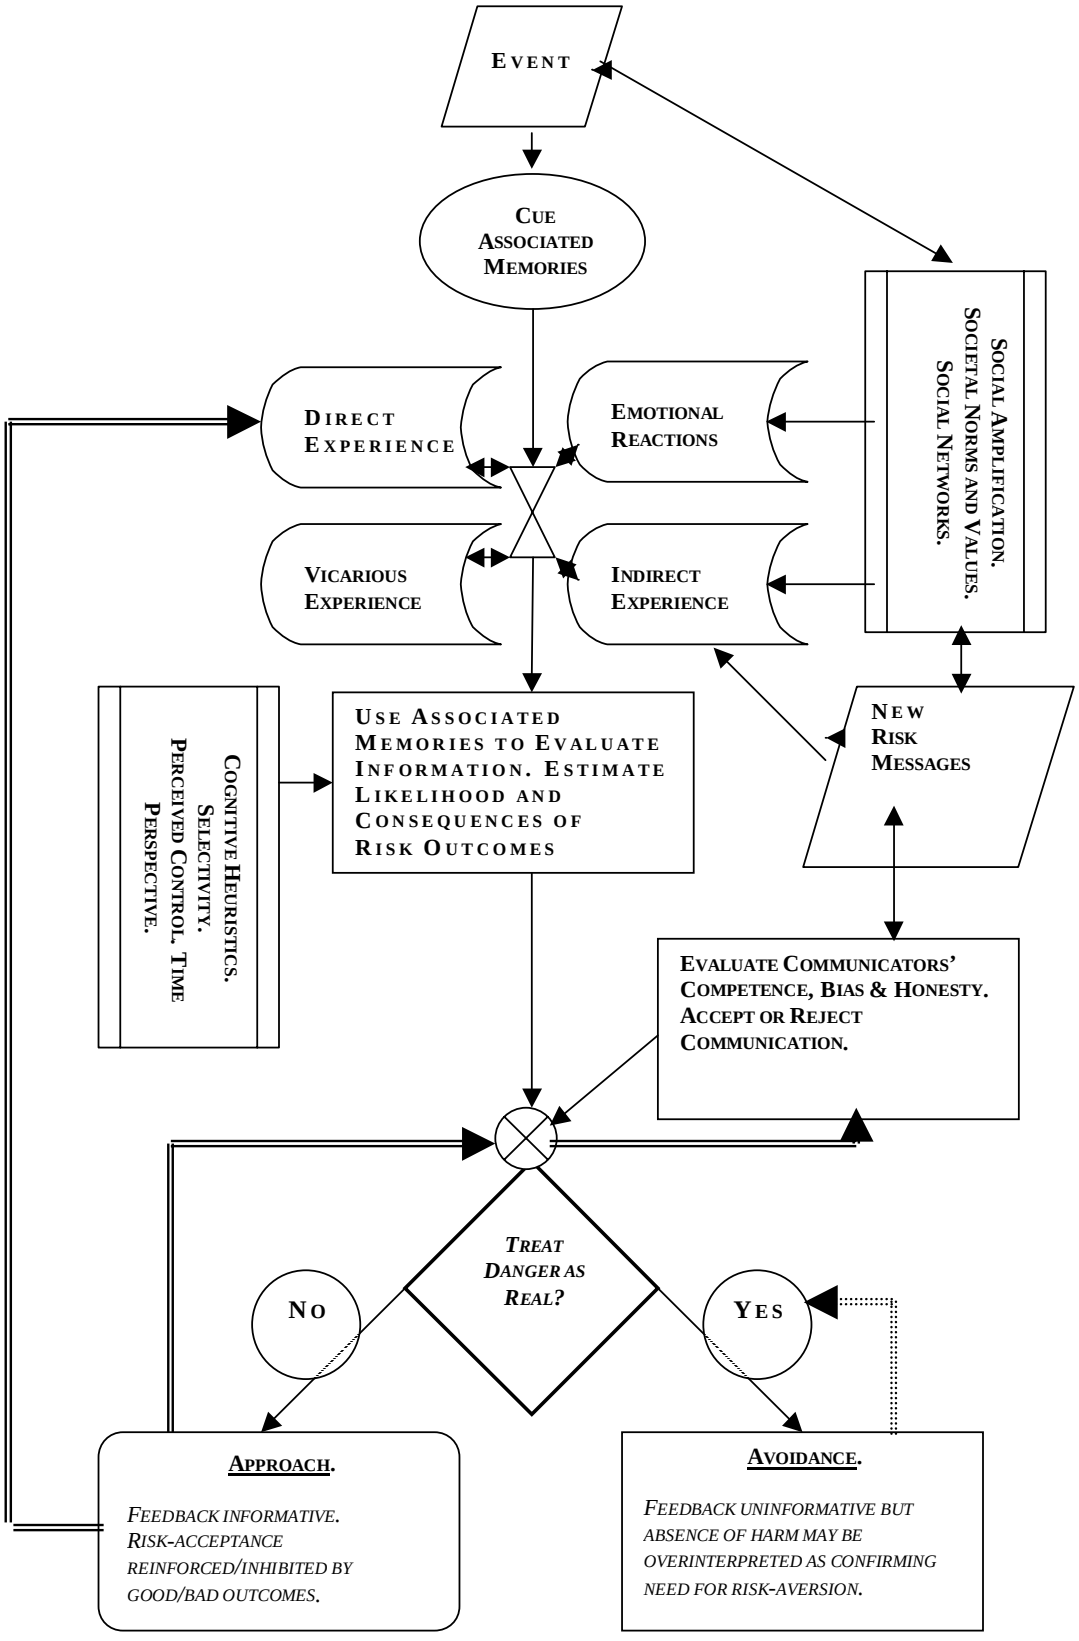
EVENT
CUE ASSOCIATED MEMORIES
DIRECT EXPERIENCE
EMOTIONAL REACTIONS
VICARIOUS EXPERIENCE
INDIRECT EXPERIENCE
SOCIAL AMPLIFICATION. SOCIETAL NORMS AND VALUES. SOCIAL NETWORKS.
COGNITIVE HEURISTICS. SELECTIVITY. PERCEIVED CONTROL. TIME PERSPECTIVE.
USE ASSOCIATED MEMORIES TO EVALUATE INFORMATION. ESTIMATE LIKELIHOOD AND CONSEQUENCES OF RISK OUTCOMES
NEW RISK MESSAGES
EVALUATE COMMUNICATORS' COMPETENCE, BIAS & HONESTY. ACCEPT OR REJECT COMMUNICATION.
TREAT DANGER AS REAL?
NO
YES
APPROACH.
FEEDBACK INFORMATIVE. RISK-ACCEPTANCE REINFORCED/INHIBITED BY GOOD/BAD OUTCOMES.
AVOIDANCE.
FEEDBACK UNINFORMATIVE BUT ABSENCE OF HARM MAY BE OVERINTERPRETED AS CONFIRMING NEED FOR RISK-AVERSION.

Part II.
Case Studies

CASE STUDY A
Planning for Disasters, with reference to the Flood and Coastal Defence Project

Two sections of the report of the Foresight Project on Flood and Coastal Defence (FCD) are particularly relevant to the issue of public perception of risk and to the themes addressed in Part one of this review. These are sections A13 'Stakeholder behaviour' and A14 'Public attitudes and expectations' in Volume 1 Appendix A ('Driver descriptions – catchment and coastal').

On stakeholder behaviour, the following remarks provide a starting point for discussion:

> "All stakeholder behaviour is fashioned by concerns that are motivated by a variety of factors ranging from true concern about flood risk to peripheral affairs, matters of process, or vested interests... Each stakeholder's actions will impinge upon others such that they can be expected to adjust their own positions. This could take the form of strategic manoeuvring or simply a response to new circumstances. Stakeholders exhibit varying interests, beliefs or values, some of which are not specifically related to flood risk itself but to other aspects of the management of those risks which may figure directly or indirectly in triggering responses"...
> "At a deeper level, theory suggests that the real problem of risk acceptance is not the substantive issue of flood risk but the wider moral questions regarding the trade offs involved in any particular decision choice and the processes by which those choices were made. The danger comes not so much from the presence of flood hazards but from the transgression of norms to which particular social groups subscribe. All of this suggests that stakeholder behaviour in respect of flood risks over the long term is not amenable to the kind of forecasting that we can apply to physical parameters of climate change."
> "There is also an argument for investigating how and why stakeholders form the opinions that they do, and how this translates into behaviour."
>
> With respect to public attitudes, the report concludes:
>
> "The most obvious ways in which public preferences are likely to influence flood risk are through public reactions to alternative decisions on flood-risk management. In particular, the acceptability of any imposed regime of flood-risk management, and its associated actual risk, will depend on a number of variables including the perceived risk and its tolerability, cost of intervention and who pays, any equity issues, any undesired consequences of interventions, alternative styles of intervention and their attributes, the process by which choices are

made, and trust in the 'system' including the people and institutions involved."

The following figure, reproduced from p. 270, corresponds closely to the signal detection paradigm described in Part 1.

	Risk is small	Risk is not small
Risk perceived as small	A. Not a driver	B. Is a driver
Risk perceived as significant	C. Is a driver	D. Is a driver

"In circumstance A, Public Attitudes and Expectations is not a driver of flood risk. In circumstance B, it is a driver – it exacerbates flood risk because no action is taken. Given the media interest in floods, this may not happen. In C and D, Public Attitudes and Expectations will act to try to reduce risk."

Building on these observations and other findings of the report, the following themes deserve particular attention.

Uncertainty of forecasting

Predicting precisely when, and with what severity, any future flooding disaster will occur is very uncertain, since any such disaster depends on the co-occurrence of different risk conditions, including stakeholder behaviour. Greater confidence can be attached to identifying where any flooding that does occur is likely to cause most damage. Hence, there are issues regarding public perception of both the overall and relative level of any risk. An important question is whether members of the public will find it easier to appreciate relative levels of risk than absolute levels. For example, they may find it easier to appreciate that flooding is most likely in low-lying and coastal areas than to put a numerical figure on such likelihoods.

Climate change

Particular issues arise with flood risk because of climate change, which causes the anticipated risk of flooding to increase, rather than vary randomly, over time. The FCD report presents projections of this increase according to different assumptions concerning economic growth and strategies for reduction of emissions. These assumptions carry different implications of costs and benefits. Furthermore, the time-frame of such projections is such that costs may need to be incurred now to yield benefits, or reduce costs, for future generations. For most people alive today, this constitutes not simply delay of gratification, but delay beyond their own lifetime.

Decision-making and behaviour

Natural disasters such as floods are not uniquely attributable to any specific kind of human activity (hence the phrase ‘act of God’) but their effects can be exacerbated or mitigated by what people do. Climate change, which in may greatly increase flood risk, is very largely related to human industrial activity and the consumption of fossil fuels. House-building in flood plains puts householders in harm’s way, whereas investment in river and coastal defence schemes can reduce such risk. The trouble is that all risk-reduction measures have costs as well as benefits, and all activities that increase risk have benefits as well as costs. Furthermore, such benefits and costs may be distributed unevenly - those who bear most costs may not at all be those who reap most benefits - raising issues of equity.

How are such costs and benefits estimated and perceived?

There are two distinct sets of questions here. One is how people estimate the costs, to themselves and others, of being a victim of a flood, and hence the perceived benefit of avoiding being flooded. The second is how they estimate the efficacy of particular risk-reduction measures, or the impact of activities that increase the risk.

The first issue is one of how direct personal experience, vicarious experience, and communications from others combine together. Other things being equal, personal experience of being a victim of flooding should greatly increase people’s perception of the likelihood and severity of such an experience. Especially if the people concerned continue living where they did before, an increased perception of risk can be regarded as a rational generalisation from their previous experience. However, this may be partly off-set by motivational factors, particularly if the individuals concerned feel emotionally and/or financially committed not to move home. Such individuals may wish to play down the level of risk (cf. Janis & Mann, 1977). They might admit that their home is at somewhat greater risk than the average, but do not want to think of it as being at high risk, like smokers who admit their risk of cancer is higher than average, but not as high “as people say”.

Direct personal experience, however, can also act to keep risk perceptions low. Since serious floods are rare, most people’s personal experience may lead then to assume that they are in little personal danger from them.

Vicarious experience, the observation of others’ experience, should generally lead to higher perceptions of risk. This is because, when floods occur, they are likely to receive graphic reporting, including television footage of flooded towns and villages, homes damaged by flood, families who have lost their homes and possessions, and even fatalities. This makes social amplification of the risk message likely. By contrast, the absence of a flood is not worth reporting even where a ‘non-event’ represents the successful efforts of those designing and implementing flood protection schemes.

Indirect experience appears to be divided into warnings (or denials) of risks of global climate change at one extreme and public consultations regarding local flood defence schemes at the other, with not a lot in between. Regarding the

former, potentially chaotic or catastrophic changes in a complex dynamic system are difficult to predict, even for experts. An important challenge is to communicate the certainty that climate change is happening along with the considerable uncertainty about what this may mean in terms of weather patterns. This very uncertainty is itself a source of risk, in that it makes protective planning more difficult, both at a personal and at a policy level. This complicates the communication of scientific opinion. If 'experts' can't agree what's going to happen, the public might think, their expertise isn't worth very much and little confidence need be placed in their opinions. A possible way out of this is to attempt to communicate that scientists do agree that weather patterns will show more extreme fluctuations even if they cannot agree on the fine detail.

In risk communications in the context of local consultations, the contrast between 'expert' scientific opinion and 'inexpert' members of the public is less clear-cut. Local residents may be less aware of the theoretical models used, for instance, to predict tidal flow, or the relationship between rainfall and rising river levels, but they can offer local knowledge, often gained over many years, on such topics as the erosion of river banks, which fields and roads appear to be the first to get inundated, or the disruption caused to traffic by a flood. Opinions of 'experts' are thus more likely to be trusted, so long as they seem compatible with residents' direct experience.

Suggestions about what to do about such problems are likely to be more contentious. The problem of individuals appreciating, or being persuaded, that their behaviour is environmentally damaging and needs to be changed resembles the situation when people judge their own health behaviour, or how good a driver they are. If adverse consequences are infrequent (as with floods) or accumulate over long periods of time (as with climate change), individuals' personal experience fails to provide the negative reinforcement, or corrective feedback, that could be an impetus for reappraising risk and changing behaviour. Especially with climate change, the impact of any individual's behaviour is imperceptible. Acknowledging that human activity has an impact, or even that climate change is happening at all, relies upon accepting the existence of a causal relationship, which is something that comes from scientific analysis and not personal experience. Individuals may feel that their memory of seasonal variation in previous years is indicative of global warming, though this kind of memory search seems to involve looking for evidence supporting an already-formed hypothesis. Being told about global warming leads one to recall cold winters from decades earlier more readily than mild ones.

If direct experience does not reveal the need for, or efficacy of, personal behavioural change, the impetus for change must come from others. Social influence can be very important here. In the case of protection from an immediate flood hazard, local communities and social networks may be important sources of help and advice, perhaps more in mitigating the effects of flooding than in preventing its occurrence. Local networks are important when problems such as the provision and use of sand-bags, evacuation, or care of elderly residents arise.

If individuals think about climate change on a global scale, they may choose to be less wasteful in their use of energy and natural resources because such behaviour is a ‘good thing’ for the environment, with beneficial effects such as a decrease the rate of global warming and associated flood risk. However, it is uncertain how much individuals are concerned with specific consequences such as the reduction of long-term flood risk, reducing the need for more roads to cope with increased car use, more power stations, or more use of landfill, or are simply adopting a more general code of thrift and good husbandry. The challenge of personally doing something to reduce future flood risk from climate change appears to be subsumed into a more general category of ‘environmentally friendly’ life-styles and norms, without people necessarily articulating a mental model of the links between energy use, climate change and flooding.

For example, in many countries including the UK there are increasing social pressures, as well as regulations and bylaws, that encourage or require citizens to send various kinds of household waste for recycling. In some countries, this almost has the flavour of a moral imperative, with neighbours quick to remind recalcitrant newcomers of their civic duty. Here, the social costs of not recycling start to outweigh the inconvenience of doing so. But there is little public debate about the benefits of recycling, at least in terms of energy saving. Recycling consumes energy, and the cost-effectiveness of such energy use may often be debatable. Alternative strategies for waste disposal, in particular incineration, may be very unpopular. Objections to incineration tend to be expressed in terms of fears of potential risks to health, although the visual impact may also be a factor.

Whereas recycling seems to be increasingly accepted as a way in which individuals can ‘do their bit’, less ritualistic but potentially more effective changes in behaviour are difficult to achieve. Car and air travel produce significant carbon emissions, but also provide benefits that cannot be foregone without changes in people’s lifestyles. Such benefits are experienced immediately and directly, which produces a bias against change. The critical element here is whether people feel they have a real choice over such aspects of their behaviour. People will use cars for pleasure and convenience, but will also invoke the ‘excuse’ that they have no choice if public transport is inadequate, or if they need their car for work. This perceived lack of personal control can inhibit behavioural change even where it is obviously desirable.

The responsibility for change does not fall merely on individuals, but also on industry and government. These bodies also influence public pronouncements, for example those by the current US administration over the risk of climate change and the need for any reduction of carbon emissions. Attempts by governments to downplay the need for changes that might have a negative short-term economic impact invite attributions of response bias. Perhaps this reflects a biased interpretation of the level of risk (could anyone really believe that carbon emissions aren’t causing climate change?), but at least seems to include a communicative response bias. The tobacco

industry's denial of any link between smoking and cancer comes to mind. All the elements are there for such pronouncements to be seen as ill-informed, self-interested and insincere – in short, not to be trusted.

If we consider local flood defence schemes, rather than policies to counteract climate change, perceived costs and benefits also have an impact on public trust. Any scheme, from raising the Thames Barrier to reinforcing river banks, involves a cost. Calculating the value for money of such schemes requires technical expertise, and one of the hardest things to communicate in a local public consultation could be that a scheme that residents believe would be desirable would produce only a marginal or short-term reduction in risk, and at an excessive cost. If limited resources allow only some homes and businesses to be offered the same high level of protection, perceptions of inequity are likely to result, especially since all tax-payers and rate-payers will be contributing to the costs. Those left out may feel their needs are disregarded and that policy-makers are biased against them. To counter such suspicions, a start could be for the consultation to involve communicating, and inviting comments on, the assumptions underlying such cost-benefit and value-for-money calculations, rather than merely presenting the results of such calculations as a fait accompli.

CASE STUDY B.
New Technology, with reference to the Cyber Trust and Crime Prevention project

Public perceptions of the risks of new technologies have been relatively widely researched, but there is still a need for theoretical integration. It cannot simply be assumed that all technologies will be represented and evaluated similarly. In particular, it is important to distinguish between new technologies that have already been brought to market, and of which the public may have some direct experience, and others which are still at the research stage so that public experience is largely indirect. Nonetheless, some general themes can be identified that are common to both stages.

Trust

Probably the most important theme is that of trust. There are several interrelated reasons for this, but most stem from a lack of relevant expertise on the part of the individual citizen. Without the requisite expert knowledge, ordinary people can't tell what makes a technological system more or less safe and may not even be able to detect any harmful consequences that it may involve, let alone attribute them to the technology itself rather than to some alternative cause. Experts need to be relied upon both to evaluate whether the technology can be safely developed in the first place, and to manage and regulate risks associated with its use. By the same token, however, if experts were always trusted, there wouldn't be a public perception problem to be addressed. Despite having little choice but to rely on experts, we don't tend to put ourselves completely in their hands and switch off our own judgemental processes.

Associations

Previous experience (direct, vicarious or indirect) can sometimes produce distrust and an elevated fear of risk, usually through generalisation from hazards or events that appear similar to the technology being considered. The degree of similarity is essentially a matter of judgement. It depends on one's mental model of how the technology works and how it could go wrong. Some (but not all) of the opposition to the building of nuclear power stations during the 1970s and 1980s could be attributed to the perception that these stations were exploiting the same technology as that used to make nuclear weapons. This led to a moral stigma for nuclear power and to fears that nuclear power stations could lead to fatalities on a massive scale. Language assists such generalisation, as with terms such as 'radiation' that extend to technologies involving electro-magnetic frequencies. Emotional reactions can also be influential, which is not to dismiss these as irrational. A number of new (especially reproductive and genetic) technologies raise ethical issues which, at least from the perspective of some belief systems, challenge traditional assumptions about the meaning of human life and our right to 'tamper with nature'. Such objections may make good copy for the press and broadcasters, and lend themselves to social amplification.

Technologies already in the market

Direct experience of a technology can lead to more positive perceptions, even if we have little understanding of how it actually works. This is clearest for technologies that have already reached the market, such as mobile phones and electronic commerce. The benefits of using them can create favourable attitudes and even forms of dependence, so that we may arrive at the point of saying that we "couldn't manage" without our mobile phone, or the possibility of purchasing goods and services over the internet. Such experience may potentially be imbalanced in that these benefits may be more obvious than possible risks to our health, wealth or privacy. It may be very difficult to detect the presence of such risks unless something bad happens – and bad happenings, as we have seen, can be statistically infrequent even in comparatively hazardous situations.

The standards by which people compare new technology vary widely, and not just for the public as opposed to experts. Individuals may rely upon memories of experiences before the technology was introduced or when it was in its early stage of introduction. For users of the technology among the general public, comparisons with the past are likely to be mostly positive. Internet connections become faster, websites are easier to find, reception on mobile phone calls becomes clearer and more reliable. In terms of the feedback model proposed in Part I, such experienced changes are likely to lead us to 'approach' or use it (use a mobile phone, make internet bookings and purchases) and, so long as nothing obvious goes wrong, our positive expectancies will be reinforced and our own use of the technology will become more frequent and established.

The notion of an 'affect heuristic' (Slovic *et al.*, 2002) suggests that beliefs in the existence of benefits should inhibit thoughts about possible dangers. This is good news for the manufacturer. But for the policy-maker charged with protecting public health, it implies that the public will be resistant to precautionary messages about the possible risks. Even so, the possibility of risk may not be discounted entirely. Although those who see most benefits may see fewer risks, the affect heuristic suggests that users will regard these risks as acceptable rather than absent. Users may differentiate between contexts in which the risks are unacceptably and avoidably large, and those where they are more tolerable and perhaps personally controllable. For example, in our own research on public attitudes towards mobile phone technology (Eiser & White, 2003), we find that greater use is associated with lower concern about health risks. Similarly, those reporting using mobile phones while driving, before the recent legislation against doing so, regarded this activity as less dangerous than others did (White, Eiser & Harris, 2004). However, mobile phone users reported concerns about possible risks to children and were rather more worried about the location of mobile phone masts and base stations. This confirms a general rule that risks that are seen to be imposed rather than voluntarily chosen or personally controllable may arouse objections that could appear out of proportion to the relative dangers involved.

Within the context of electronic commerce and cybertrust, differences in comparison standards are likely to lead to very different perceptions of risk, reflecting individuals' understanding of the workings of electronic systems and hence their vulnerability. It is a truism that nothing in life is risk-free, and this certainly applies to electronic commerce.

However, from the perspective of the non-expert public, the critical issue is not whether any risk is present, but whether it is perceived as being more serious than what went before. In the UK, we are currently being told that credit card transactions in shops will soon be requiring us to type in a PIN rather than sign a voucher, as already happens more frequently elsewhere. The presumption is that this will be, and be seen to be, an improvement in security. Under the old system, if my credit card was stolen, how easily could someone forge my signature (reasonably easily) and how easily might such a forgery survive the cursory inspection (if any) given it by many sales staff (again, reasonably easily)? Or how easily could someone – a dishonest salesperson for instance – make a copy of my signature and card details from a voucher or receipt? Again, this is something well within the bounds of imaginability. These, however, are familiar risks. The risks associated with reliance on purely electronic identification, however, I find more difficult to assess. Could anyone discover my PIN by looking over my shoulder? Imagining this is easy enough. The location of many cash machines on busy streets already compromises security in this way. But what is a mystery to me is what it would take for the electronic storage of my account details, card number and PIN to be penetrated by someone with fraudulent intent, or to be made effectively impenetrable. Because I have only a primitive mental model of electronic transaction systems, I have no way of knowing directly how secure they are, beyond the fact that I have – touch wood – had no disasters so far. My own experience coincides with the assurances I receive from my credit card company.

But for the experts charged with making such systems as secure as possible, this isn't quite enough. They need to stay ahead of the game. They need to be able to identify possible weaknesses in the systems that are vulnerable to error or fraudulent attack. If they can make the systems even more secure against such events, they have something extra to offer corporate and individual customers in a competitive market. The pay-offs therefore favour innovation. And of course, the game is being played not merely with commercial competitors, but against criminals who are also attempting to identify and exploit weaknesses in the current systems and for whom, likewise, profits depend on innovation.

Risk-taking in the business sense of the acceptance of innovation thus involves here a heightened awareness of how an apparently secure system could be attacked, and risk-averse decision-making that involves treating possible dangers as real. In other words, they may compare the security of current systems not merely with the paper transactions of a few years ago, but with something even more foolproof and fraud-proof than we have now. A system that feels secure to ordinary members of the public may be interpreted as imperfectly secure by experts. The experts may also rely on vicarious

experience in terms of observation and hearsay concerning companies which failed to adopt new defences and lost out as a consequence.

Feedback is likely to reinforce this combination of risk-averse perceptions and innovation, so long as the costs of research and development of new security processes do not too greatly exceed the estimated costs of system failure. Companies which are aware of the possible weaknesses in their system are unlikely to accept the risk of not investing in improved security within the limits of what they can afford. Feedback will tend to reinforce this risk-averse orientation through offering decision-makers little information about what might have happened if the extra security processes had not been developed or adopted. In other words, true-positive and false-positive (false-alarm) responses to the presence of risk cannot be easily discriminated.

Of most of this, ordinary members of the public – unless they have personally been defrauded – will for the most part be blissfully unaware. Probably the main issue for users will be the acceptability of new security procedures that involve extra cost (which might be hidden within credit card charges, perhaps), inconvenience (needing to remember and protect the secrecy of a PIN) and general concerns about personal and financial data protection. The first two of these may be balanced against possible perceived benefits, but if these benefits are not very obvious, customers may resent what they see as the imposition of unnecessary expense and hassle. The proposed introduction of national identity cards to the UK may be seen to pose a similar problem. It is difficult for citizens to make balanced estimates of any benefits to society as a whole against personal inconvenience and feared loss of privacy. Concerns about data protection probably depend largely on trust. Since most of us will have rather little knowledge of the inner workings of such technologies or of the risks that might be increased or decreased by new procedures, much will depend on our judgements of the competence and honesty of companies, banks, police forces or governments. Judgements of the technology turn into judgements about people.

A distinctive feature of cybertrust concerns, however, is that we don't have to just consider those who are involved in researching, developing or regulating the technology. We also need to take account of the possibility that there could be someone deliberately trying to damage our wealth. Our judgements of risk could be affected by our thoughts concerning how skilful, motivated and numerous the criminals are who might attempt to penetrate such security. It is an open question how often, if at all, ordinary people engage in this form of thinking. However the frequency with which internet and email systems are attacked by computer viruses could be interpreted as evidence that there are enough people out there with the skills and motivation to cause nuisance and damage. This suggests that the industry should not be complacent about its perceived competence in the eyes of the public. Experts they may be, but with more and more people acquiring computer skills, are they as expert as they need to be?

Technologies further from market

With new technologies that are still at earlier stages of research, public perceptions of risk cannot be based on actual experience of benefits and costs. Even so, public perceptions can be very influential in shaping policies that facilitate or inhibit the research and development required to bring such technologies to market. Where might such perceptions come from? The main influences would seem to be:

a) the generalisation of beliefs from more familiar technologies;
b) the generalisation of emotional reactions to these technologies;
c) opinions expressed by other people who claim to know more about the technology in question; this last influence being moderated by
d) the degree of trust in these other people (or people of the same class, e.g. 'scientists' or 'environmentalists').

With regard to a and b, any comparison between unknown and better known technologies will inevitably involve some oversimplification. What is critical will be the apparent similarities and differences that are treated as relevant. This is likely to be a highly subjective matter, with scope for cognitive heuristics to exert an influence for better or worse. The question faced by any person judging the technology is essentially: What kind of thing is this? The answer can involve common names or labels, but they can differ considerably in terms of precision and implied evaluation. Thus 'life sciences' can sound like a Good Thing, whereas 'genetic modification' or 'genetic engineering' seems more worrying. 'Frankenstein foods' is deliberately alarming. It is understandable that scientists and technologists prefer to use more neutral terms that mean little to the uninitiated. But obscurantism cannot remove the need to explain oneself and can give rise to the suspicion that there is something to hide. If scientists do not accept the responsibility of communicating what they are doing, others will do it for them, and not always with friendly intent.

With respect to factors c and d, sociologists (e.g. Beck, 1992) have stressed cultural shifts that have led to a supposed disenchantment with modernism, science in general, and scientists in particular. From a more psychological perspective, however, there is still a need to explain how such cultural shifts occur at the level of the individual. Within any society at any given time, some individuals will be more trusting and others more distrusting of scientists, and of their critics. Such individual differences in attitude need to be accounted for. The framework outlined in this review suggests one possible starting point, which is to look at the learnt associations and experiences that lead each individual to trust some sources and distrust others. If a given industry, or 'science' generally, is seen to have a poor track record in terms of identifying possible dangers from a technology, it is likely to be distrusted on grounds of a lack of competence or *discrimination ability*. If scientists are seen as being in the pay of an industry intent on promoting a technology, their reassurances are likely to be distrusted as reflecting a *response bias* in their willingness to agree that ambiguous information suggests the presence of danger, and to communicate such concerns if they have any.

Another difficulty with accounts based on theories of broad cultural trends is that they could suggest (not least with the benefit of hindsight) a kind of historical inevitability in the shifts of opinion that occur on different issues. In practice, public opinion and opinion change is a dynamic system. It can be unstable and vulnerable to chance events and interventions. It is likely to be especially unstable and unpredictable with respect to new phenomena where prior attitudes are based on the loose use of vague heuristics and are reinforced by little direct experience. Under such circumstances, people's evaluations will be relatively weak and poorly informed. A single press article, TV documentary or popular book can have a disproportionate influence, both immediately and when it is taken up by the media and other agents of social amplification.

One example of this is the imagination of possible futures involving the exploitation of nanotechnology, initiated particularly by Drexler (1992, 2001). On the one hand, we are offered the prospect of minuscule robots or 'nanobots' that can hunt down cancer cells, but on the other hand, we are menaced with the prospect of artificial bacteria that, if their ability to self-replicate is not controlled, "could spread like blowing pollen, replicate swiftly, and reduce the biosphere to dust in a matter of days," the "gray goo" scenario (Wood, Jones & Geldart, 2003).

So what is the way forward? It must involve creating the context for informed debate and decision-making that gets beyond labels and slogans, considers the details of possible benefits and costs, and acknowledges the uncertainty of its predictions.

All this requires a shift to less reliance on heuristic and selective processing and more attention to detailed information. This transition may not be straightforward. There is, first of all, the question of whether the right kind of detailed information will be sought and provided. Quite often debates focus on the potential costs and benefits of adopting a technology, but attend less to the costs and benefits of not adopting it (misgivings about the MMR vaccine are a possible example). Questions of necessity and opportunity costs are often not examined as systematically as they should be. Next, the detailed information will often be complicated and so the audience, whether public or expert, will need an appropriate model in terms of which to interpret it. For example, the perceived 'unnaturalness' of combining genetic material from different species of animals, or even between animals and plants, might be lessened by reconceptualising DNA as a common 'digital code' for all forms of life. Finally, considering and understanding such issues in detail involves a lot of time and cognitive effort. As with all forms of information search, will the gain in knowledge justify how hard one has to work to get it?

CASE STUDY C
Drugs and Behaviour, with reference to the Brain Science, Addiction and Drugs project

Issues about the perception of risk in relation to drugs and behaviour fall under two main headings. How do individuals who use drugs perceive the risks involved? And what risks to themselves personally and to society in general do people believe are associated with the availability and use of drugs by other people?

Drug use and decision-making

One of the most misleading assumptions about drugs and drug use is that drug users are fundamentally different in the way they think and make choices about their lives from 'the rest of us'. Drug users are ordinary people, and there is no evidence that they are less intelligent or capable of rational thought than anyone else. On the other hand, some forms of drug use incontrovertibly wreck people's lives. So why do we do it? A starting point is to consider two distinct questions. First, what are drugs? And second, why and how do certain forms of drug use come to be seen as a problem?

What are drugs?

A very broad definition of a drug is any substance that is deliberately taken into the body for reasons other than its nutritional value. Even this definition, however, can lead us into trouble. 'Nutritional value' is neither an uncontroversial concept nor something that is immediately and accurately recognisable in all cases. The taste preferences that lead us to enjoy sweet or fatty foods may guide us efficiently towards rich sources of calories and proteins. For hunter-gatherers in a world of scarcity and danger, such taste cues can indeed be signs of 'nutritional value'. But in a world of plenty and longer life-expectancy, these cues can motivate 'unhealthy' eating with consequences such as obesity and heart disease. All of us eat for reasons of pleasure or comfort, rather than simply to stave off starvation.

The point of this analogy is not to claim that drugs are foods or foods are drugs (although compulsive eating disorders share much in common with other addictive behaviours), but to underline the fact that our eating behaviour is guided by cues that signal the availability of rewards, broadly defined, some of which will be unconditioned or inbuilt, and some learnt through experience and association. We now know vastly more than we did a generation ago about the brain mechanisms that underlie the recognition and experience of such rewards. In extremely crude terms, feelings of pleasure, pain, etc. are associated with the activation or inhibition of specific areas of the brain that are sensitive to specific chemical 'signals' produced within our bodies in response to particular events such as physical pain. Again extremely crudely, it is possible to 'cheat' this system by stimulating the brain with chemicals produced artificially outside our bodies that mimic these signals. The fact that tobacco and coca plants and white poppies provide the raw material for drugs

of abuse (whereas cabbages and cornflowers do not) is essentially a fluke of nature arising from a coincidental similarity between the chemical properties of these plants and particular neurotransmitters.

Identifying the pharmacological properties of particular drugs and the brain systems that are sensitive to them helps inform, but does not settle, broader questions of risk perception. Some drugs are illegal, others are not. Some are life-threatening to users, others, as far as we know, are not. Some seriously impede cognitive and motor functions to the point that their users constitute a threat to other people's lives, and again, others do not. Some are supplied to users through criminal networks, others through blue-chip city companies. Such classifications barely cohere with each other, let alone with any pharmacological taxonomy.

What is addiction?

This question is one where evidence from brain sciences is helpful, but which raises issues that require examination at different levels of explanation. Traditionally, addiction has been seen to be characterised by two processes believed to increase over time with repeated exposure to the drug. They are the development of tolerance, whereby the habitual user requires increased doses to obtain the same pharmacological 'benefit'; and the development of withdrawal effects when the user is denied access to the drug. This is assumed to set up a vicious cycle of ever more intense striving to obtain the drug for ever smaller positive reinforcement (and the avoidance of ever more severe negative reinforcement by withdrawal). Thus, although the underlying process is a change in the user's physiological reactions to the drug, a large part of the evidence that such a change is occurring is behavioural.

The intensity of withdrawal effects appears to be an imperfect predictor of the difficulty of abstaining or quitting, both when one looks at differences between individuals using a specific drug and when comparing different drugs. Cigarette smoking appears to be extremely addictive in terms of the rapidity with which the habit is established and the difficulty of quitting, but the effects of withdrawal – while inconvenient and uncomfortable – are far milder than for opiates or high levels of alcohol. What is not always so mild is the craving that motivates current smokers to light up when they have been without nicotine for a while. It can trigger relapse among those who are trying to quit, as well as among those undergoing treatment programmes for the abuse of alcohol and other drugs. Such a craving can still occur months after the last exposure to the relevant drug (for example, well after the individual has been 'detoxified') and hence must be distinguished from withdrawal effects which involve immediate physiological reactions to deprivation. Craving is a form of acquired motivation, something learnt, and because it is typically triggered by situational cues associated with previous use of the drug, can be viewed as dependent on memory associations.

Drug-use as risky decision-making

The decisions made by drug users conform to the framework outlined in Part I of this review. Drug use is triggered by cues from the environment and from the user's own body. These cues prompt both memory associations and potentially very powerful emotional reactions. Changes in mood, or physical sensations, can be caused by consumption of, or abstention from, drugs. The amount of deliberation surrounding such decisions varies from unreflective automaticity to self-conscious (and perhaps guilt-ridden) introspection, but this is not a distinguishing feature of addictive behaviour. All habitual behaviours can become more or less automatised, but still be guided by acquired expectancies. These expectancies can be based largely on direct experience, but vicarious experience as well as socially communicated expectancies can be influential too, for instance in influencing the interpretation of otherwise ambiguous sensory changes.

These expectancies then feed forward to the decision on whether to take the drug or not, or whether to take it at this time. This can be viewed in two ways. On the one hand, taking the drug might be seen as a form of approach or exploratory behaviour. This analogy may hold better in the context of recreational use or initiation rather than maintenance of any habit. Here potential users who take the drug in the expectation that it will make them feel good will have that expectation confirmed or contradicted by their own direct experience, whereas those who avoid such use may become more practised and confident in their ability to 'say no' but will not learn through direct experience how the drug might have made them feel.
On the other hand, taking the drug may be seen as a form of avoidance behaviour, by avoiding the discomfort or distress of anticipated withdrawal (Baker et al., 2004). This is more applicable to established habitual use and dependence. Note, however, that it is the anticipation of withdrawal, rather than withdrawal itself, that is the motivator here. If someone believes that they would feel awful or couldn't cope without the drug, and therefore have another cigarette, drink, fix or whatever, they will be avoiding feedback on whether abstinence would be quite as bad as they suspect. From this perspective, drug use by an established user can be viewed as risk-averse behaviour. This may seem a counter-intuitive suggestion because we, from our morally superior position as spectators, know about the health risks the user is incurring. But very often, the user will know about these too (although they may still downplay the seriousness of such risks; see e.g. Eiser, Reicher & Podpadec, 1995b).

Two critical factors, however, may lead the perceived risk of withdrawal to outweigh these health risks. One is the time perspective. The health costs are long term. The benefits of a high and the costs of withdrawal are immediate. The other factor, related to this, is the perception of control. Many (but by no means all) drug users may simply feel unable to abstain for any length of time. In other words, they may label themselves as 'addicts' and this may become self-fulfilling (Eiser, Sutton & Wober, 1978; Eiser & Gossop, 1979). What they can control, however, is their own immediate mood state through administration of the drug.

Perceived risks to others

Drug use causes, and is perceived to cause, a variety of societal problems that affect people other than the users. There is a danger of hypocrisy in demonising drug-users as evil or sick people set apart from the rest of society. In fact, legal drugs cause major costs to society. Millions of 'ordinary' citizens enjoy tobacco, but it is one of the major causes of death and disease in the western (and increasingly the developing) world. The main victims are smokers themselves, but risks are also displaced onto others through passive smoking, smoking by pregnant women, and fires (to say nothing of the burden on health services). Alcohol is not far behind, with its major contributions to road traffic accidents, injuries and death. It plays a part in statistics on sexually transmitted diseases and unwanted pregnancy and is strongly implicated in domestic violence.

This is not to dismiss the societal cost of illicit drug use. Drug-users (particularly users of Class A drugs such as heroin and crack cocaine) may resort to acquisitive crime or to prostitution to fund their habit. The growth of local and international networks to supply illicit drugs also raises the general level of criminality, including the use of firearms. When this activity is focussed in particular neighbourhoods, the result is a general decline in social capital.

My focus here is on the perception of risk, rather than risk itself. What might lead members of the public to regard certain drug-related risks as being more or less serious? Many members of the 'public' are, or have been, 'users' too, at least of certain drugs. They will have direct experience of some of the relevant consequences. The effect that this will have will depend both on the drug in question and the individual concerned, but may often lead to the risks being seen as smaller, because harmful consequences that are all too apparent at the level of the population may be relatively infrequent in an individual's personal experience. As with dangerous driving, many people will get away with it much of the time. The vast majority of lung cancer deaths are attributable to smoking, but most smokers don't get cancer and if they do, the symptoms will not appear for many years.

Vicarious experience is likely to act in the other direction, that is, lead to greater perception of risk. It is a cliché that many people have an elderly relative who's smoked two packs a day for sixty years and not got cancer. But most people have a good chance of knowing someone who has suffered harm from drugs, perhaps someone close to them. Bad outcomes may be easier to recall or imagine and may be seen as 'typical' consequences, despite still being relatively infrequent.

Indirect experience is likely to be somewhat more complex in its effects. Testimonials from reformed addicts recounting the horrors of their previous lives are staple fare in many drug education and cessation programmes, of which Alcoholics Anonymous is perhaps the most widely known. The highly personalised flavour of many such programmes, and of press reports of celebrities' 'battles with drugs' (or booze, gambling, etc.) makes such information almost like vicarious experience in its impact. However, not everyone is persuaded by (or keen to listen to) such stories, especially when

they seem at variance with their personal experience. Another difficulty concerns the selectivity of attention by more or less everyone involved, including the media, politicians and medical researchers (at any point in time) to certain kinds of ‘drug problems’ rather than others.

CONCLUDING REMARKS

Risk is all around us, in everything we do and in everything we experience. Risk is real, but it is not a 'thing' to be perceived or misperceived. In this review I have attempted to shift the debate away from a focus on how people make mistakes in their predictions of future events, or their calculations of costs and benefits, towards a discussion of more general decision-making processes. Of course we all make such mistakes. But these mistakes do not prove that we are stupid, still less that some sections of society (e.g. 'the public') are more stupid than others (e.g. 'the experts'), though clearly different people can have widely different levels of knowledge and experience on any given issue. It is rather the other way around. The kinds of mistakes we can make actually show how clever we are – clever at coping somehow with a highly complex environment and making decisions on the basis of uncertain evidence.

Risk is largely about uncertainty, and uncertainty is reduced through learning and experience. But what we learn and experience is steered by our perceptions of risk. We learn more about objects, people and activities we believe to be 'safe' than about those we believe to be 'dangerous', because we are generally more likely to avoid experiencing the latter. Such risk-aversion is a good strategy for survival and relative prosperity under many circumstances, even if it leaves us with areas of ignorance or even uncorrected prejudice. At the same time, our experience may provide only ambiguous or delayed feedback about our estimates of risk, so we may feel we are safe when we are actually in danger. Again, this does not show that we are stupid, only that the information we have available is often difficult to interpret.

But risk is not just about uncertainty or probability. It is also very much about value – about the desirability and undesirability of different outcomes. To the extent that we want and cherish different things, or find different threats more frightening or more repugnant, we will end up with different preferences, however accurate our prediction of future contingencies. Understanding such differences in people's values is an essential step to understanding human behaviour.

References.

Alloy, L.B., Abramson, L.Y. & Francis, E.L. (1999). Do negative cognitive styles confer vulnerability to depression? *Current Directions in Psychological Science, 8,* 128-132.

Asch, S.E. (1956). Studies of independence and conformity: A minority of one against a unanimous majority. *Psychological Monographs, 70,* (9, Whole No. 416).

Baker, T.B., Piper, M.E., Fiore, M.C., McCarthy, D.E. & Majeskie, M.R. (2004). Addiction motivation reformulated: An affective processing model of negative reinforcement. *Psychological Review, 111,* 33-51.

Bandura, A. (1977). Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review, 84,* 191-215.

Bargh, J.A. (1997). The automaticity of everyday life. In R.S. Wyer, Jr. (Ed.) *The automaticity of everyday life: Advances in Social Cognition, Vol. 10* (pp. 1-62). Mahwah, NJ: Erlbaum.

Barnes, M. (1990). *The Hinkley Point public inquiries.* London: Her Majesty's Stationery Office.

Beck, U. (1992). *Risk society: towards a new modernity*. London: Sage.

Bower, G.H. (1981). Mood and memory. *American Psychologist, 36,* 129-148.

Chaiken, S. (1980). Heuristic versus systematic information processing and the use of source versus message cues in persuasion. *Journal of Personality and Social Psychology, 39,* 752-766.

Damasio, A. (1994). *Descartes' error: Emotion, reason and the human brain.* New York: Avon.

Denrell, J. (2003). Vicarious learning, under-sampling of failure, and the myths of management. *Organization Science, 14,* 227-243.

Denrell, J. & March, J.G. (2001). Adaptation as information restriction: The hot stove effect. *Organization Science, 12,* 523-538.

Detweiler, J.B., Bedell, B.T., Salovey, P., Pronin, E. & Rothman, A.J. (1999). Message framing and sunscreen use: Gain-framed messages motivate beach-goers. *Health Psychology, 18,* 189-196.

Devos-Combey, L. & Salovey, P. (2002). Applying persuasion strategies to alter HIV-related thoughts and behavior. *Review of General Psychology, 6,* 287-304.

Drexler, K.E. (1992). *Engines of Creation: The coming era of nanotechnology*, Oxford: Oxford University Press.

Drexler, K.E. (2001). Machine-Phase Nanotechnology, *Scientific American*, 16 September.

Dweck, C.S. & Goetz, T.E. (1978). Attributions and learned helplessness. In J.J. Harvey. W. Ickes & R.F. Kidd (eds.), *New directions in attribution research,* Vol. 2. (pp. 157-179). Hillsdale, NJ: Erlbaum.

Eiser, J.R. (1990). *Social judgment.* Buckingham: Open University Press.

Eiser, J.R. (1994). *Attitudes, chaos and the connectionist mind.* Oxford, Blackwell.

Eiser, J.R., Claessen, M.J.A. & Loose, J.J. (1998). Attitudes, beliefs and other minds: Shared representations in self-organizing systems. In S.J. Read and L.C. Miller (Eds.) *Connectionist models of social reasoning and social behavior.* Mahwah, NJ: Erlbaum (pp.313-354).

Eiser, J.R., Fazio, R.H., Stafford, T. & Prescott, T.J. (2003). Connectionist simulation of attitude learning: Asymmetries in the acquisition of positive and negative evaluations. *Personality and Social Psychology Bulletin, 29,* 1221-1235.

Eiser, J.R. & Gossop, M.R. (1979). 'Hooked' or 'sick': Addicts' perceptions of their addiction. *Addictive Behaviors, 4,* 185- 191.

Eiser, J.R., Reicher, S.D, & Podpadec, T.J. (1995a). Global changes and local accidents: Consistency in attributions for environmental effects. *Journal of Applied Social Psychology, 25*, 1518-1529.

Eiser, J.R., Reicher, S.D & Podpadec, T.J. (1995b). Smokers' and non-smokers' estimates of their personal risk of cancer and of the incremental risk attributable to cigarette smoking. *Addiction Research, 3*, 221-229.

Eiser, J.R., Sutton, S.R. & Wober, M. (1978). "Consonant" and "dissonant" smokers and the self-attribution of addiction. *Addictive Behaviors, 16,* 99-106.

Eiser, J.R., van der Pligt, J. & Spears, R. (1995). *Nuclear neighbourhoods: Community responses to reactor siting.* Exeter: University of Exeter Press.

Eiser, J.R. & White, M.P. (2003). *Public perceptions of the health and safety risks of mobile phones*: Report to the Health and Safety Executive.

Fazio, R.H. (1990). Multiple processes by which attitudes guide behavior. The MODE model as an integrative framework. In M.P. Zanna (Ed.), *Advances in experimental social psychology,* (Vol.23, pp 75-109). San Diego, CA: Academic Press.

Fazio, R.H. (2001). On the automatic evaluation of associated evaluations: An overview. *Cognition and Emotion, 15,* 115-141.

Fazio, R.H., Eiser, J.R. & Shook, N.J. (2004). Attitude formation through exploration: Valence asymmetries. *Journal of Personality and Social Psychology.87,* 293-311.

Fischhoff, B. (1983). Predicting frames. *Journal of Experimental Psychology: Learning, Memory and Cognition, 9,* 103-116.

Fischhoff, B., Slovic, P. & Lichtenstein, S. (1978). Fault trees: Sensitivity of estimated failure probabilities to problem representation. *Journal of Experimental Psychology: Human Perception and Performance, 4,* 330-344.

Gigerenzer, G. & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction. *Psychological Review, 102,* 684-704.

Gilovich, T. Griffin, D. & Kahneman, D. (Eds.) (2002). *Heuristics and biases: The psychology of intuitive judgment.* New York:Cambridge University Press.

Heider, F. (1946). Attitudes and cognitive organization. *Journal of Psychology, 21,* 107-112.

Heider, F. (1958). *The psychology of interpersonal relations.* New York: Wiley.

Hogarth, R.M. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin, 90,* 197-217.

Janis, I.L. (1972). *Victims of groupthink.* Boston, MA: Houghton Mifflin.

Janis, I.L. & Mann, L. (1977). *Decision making: A psychological analysis of conflict, choice, and commitment.* New York: Free Press.

Johnson, E.J., Hershey, J., Meszaros, J. & Kunreuther, H. (1993). Framing,

probability distortions, and insurance decisions. *Journal of Risk and Uncertainty, 7,* 35-51.

Kahneman, D. & Tversky, A. (1971). Belief in the law of small numbers. *Psychological Bulletin, 76,* 105-110.

Kahneman, D. & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263-291

Kahneman, D. & Tversky, A. (1984). Choices, values, and frames. *American Psychologist, 39,* 341-350.

Kasperson, R. E., Jhaveri, N. and Kasperson, J. X. (2001). Stigma, places, and the social amplification of risk: toward a framework of analysis, in J. Flynn, P. Slovic, and H. Kunreuther (eds.), *Risk, Media and Stigma: Understanding Public Challenges to Modern Science and Technology*. London: Earthscan, pp. 9-27.

Lange, R. & Fishbein, M. (1983). Effects of category differences on belief change and agreement with the source of a persuasive communication. *Journal of Personality and Social Psychology, 44,* 933-941.

Langer, E.J. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32,* 311-328.

Lerner, M.J. (1980). *The belief in a just world: A fundamental delusion.* New York: Plenum.

Lopes, L.L. (1987). Between hope and fear: The psychology of risk. In L. Berkowitz (Ed.) *Advances in experimental social psychology, Vol. 20* (pp. 255-295). San Diego, CA: Academic Press.

March, J.G. (1996). Learning to be risk averse. *Psychological Review, 103,* 309-319.

Mischel, W. (1973). Towards a cognitive social learning reconceptualization of personality. *Psychological Review, 80,* 252-283.

Mischel, W. (1974). Processes in delay of gratification. In L. Berkowitz (Ed.). *Advances in experimental social psychology,* vol. 7. New York: Academic Press.

Mischel, W. & Mischel, H.N. (1983). Development of children's knowledge of self-control strategies. *Child Development, 54,* 603-619.

Mischel, W. & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review,* 102, 246-268.

Newcomb, T.M. (1953). An approach to the study of communicative acts. *Psychological Review, 60,* 393-404.

Nisbett, R.E. & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment.* Englewood Cliffs, NJ: Prentice-Hall.

Petty, R.E. & Cacioppo, J.T. (1986). The elaboration likelihood model of persuasion. In L. Berkowitz (Ed.). *Advances in experimental social psychology* (Vol. 19, pp. 123-205). San Diego, CA: Academic Press.

Petty, R.E. & Cacioppo, J.T. (1990). Involvement and persuasion: Tradition versus integration. *Psychological Bulletin, 107,* 367-374.

Petty, R.E. & Krosnick, J.A., Eds. (1995). *Attitude strength: Antecedents and consequences.* Mahwah, NJ: Lawrence Erlbaum.

Pidgeon, N., Kasperson, R. E. and Slovic, P. (Eds.) (2003). *The Social Amplification of Risk*. Cambridge: Cambridge University Press.

Raven, B.H. (1974). The Nixon group. *Journal of Social Issues, 30,* 297-320.

Ross, L. (1977) The intuitive psychologist and his shortcomings: Distortions in the attribution process. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10). New York: Academic Press. pp.174-221.

Ross, L., Greene, D. & House, P. (1977). The 'false consensus effect': An egocentric bias in social perception and the attribution process. *Journal of Experimental Social Psychology, 13,* 279-301.

Rothman, A.J. & Salovey, P. (1997). Shaping perceptions to motivate healthy behavior: The role of message framing. *Psychological Bulletin, 121,* 3-19.

Rothman, A.J., Martino, S.C., Bedell, B.T., Detweiler, J.B. & Salovey, P (1999). The systematic use of gain- and loss-frame messages on interest in and use of different types of health behavior. *Personality and Social Psychology Bulletin, 25,* 1355-1369.

Rotter, J.B. (1966). Generalized expectancies for internal versus external control of reinforcement. *Psychological Monographs, 80,* (Whole No. 609).

Rozin, P. & Royzman, E.B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review, 5,* 296-320.

Seeman, M. (1971). The urban alienations: Some dubious theses from Marx to Marcuse. *Journal of Personality and Social Psychology, 19,* 135-143.

Seligman, M.E.P. (1975). *Helplessness.* San Francisco, CA: Freeman.

Slovic, P., Finucane, M., Peters, E. & MacGregor, D.G. (2002). The affect heuristic. In T. Gilovich, D. Griffin & D. Kahneman (Eds.) *Heuristics and biases: The psychology of intuitive judgment* (pp. 397-420). New York: Cambridge University Press.

Stewart, I. (1989). *Does God play dice? The mathematics of chaos.* Cambridge, MA. Blackwell.

Strategy Unit (2002). *Risk: Improving government's capability to handle risk and uncertainty.* London: Strategy Unit, Cabinet Office; www.strategy.gov.uk.

Sutton, R.S. & Barto, A.G. (1998). *Reinforcement learning: An introduction.* Cambridge, MA: The MIT Press.

Swets, J.A. (1973). The receiver operating characteristic in psychology. *Science, 182,* 990-1000.

Tversky, A. & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185,* 1124-1131.

Tversky, A. & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science, 211,* 453-458.

Tversky, A. & Koehler, D.J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review, 101,* 547-567.

van der Maas, H.L.J., Kolstein, R. & van der Pligt, J. (2003). Sudden transitions in attitudes. *Sociological Methods and Resarch, 32,* 125-152.

van der Pligt, J., Eiser, J.R. & Spears, R. (1986). Construction of a nuclear power station in one's locality: Attitude and salience. *Basic and Applied Social Psychology, 7,* 1-15.

White, M.P., Eiser, J.R. & Harris, P. (2004). Risk perceptions of mobile phone use while driving. *Risk Analysis,*

Wood, S.J., Jones, R. & Geldart, A. (2003). *The social and economic challenges of nanotechnology: A report to the ESRC.*

Zajonc, R.B. (1980). Feeling and thinking: Preferences need no inference. *American Psychologist*, 35, 151-175.

Zeeman, E.C. (1976). Catastrophe theory. *Scientific American, 234,* 65-83.