David Hardman

[redacted]

Judgment and Decision Making

Psychological Perspectives

[redacted]

[redacted]

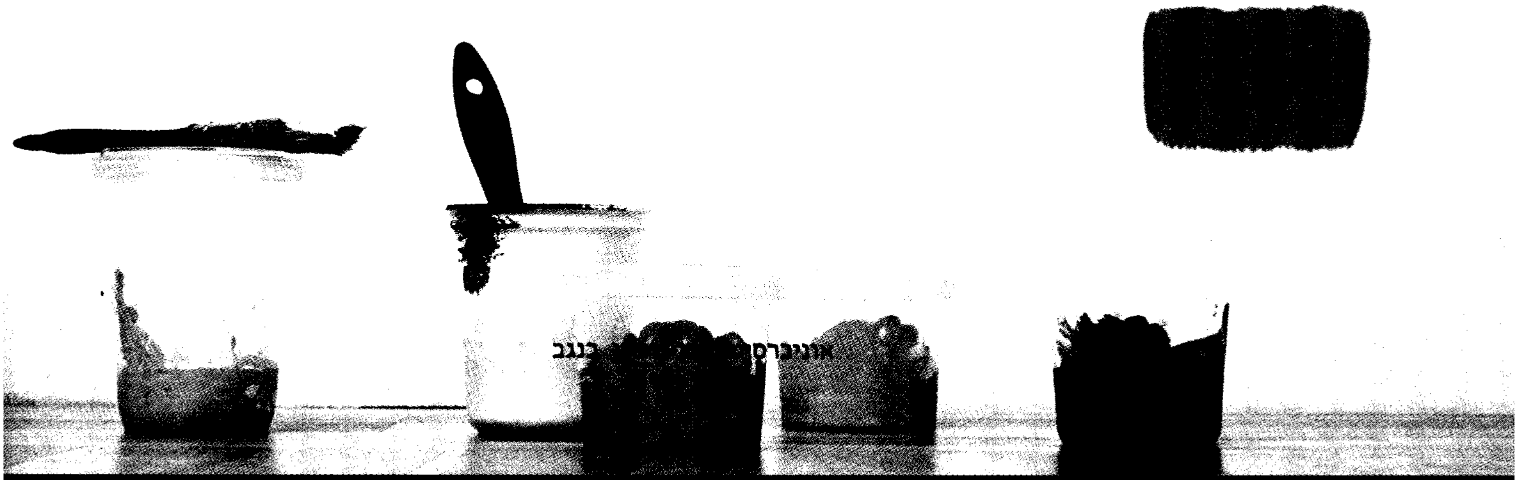

BPS TEXTBOOKS IN PSYCHOLOGY
BPS Blackwell

7 Decision Making under Risk and Uncertainty

KEY TERMS

Allais paradox
ambiguity aversion
certainty effect
coefficient of variation
cumulative prospect theory
decision by sampling
decision field theory
disappointment
expected utility
expected value
framing effects
functional magnetic resonance imaging (fMRI)
ignorance aversion
loss aversion
possibility effect
priority heuristic
probability weighting function
prospect theory
reflection effect
risk sensitivity
St Petersburg paradox
source preference
subadditivity
two-stage model
value weighting function
Weber's Law

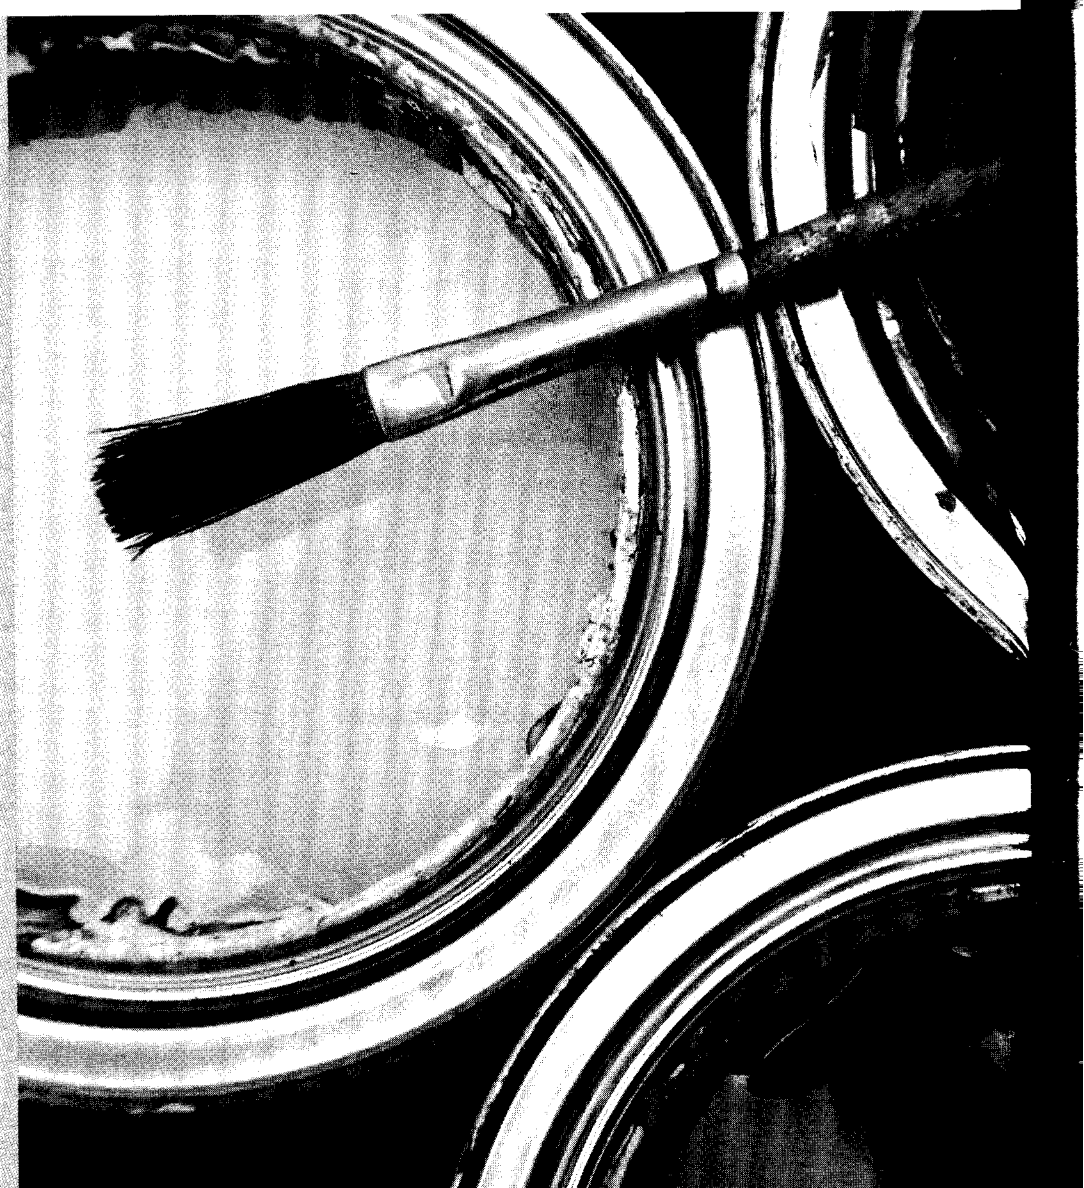

CHAPTER OUTLINE

Before reading any further, please give your responses to the following problems (for numerical outcomes, just assume that the units are the major units of your own national currency).

Problem 1
Which would you prefer:
A. A certain gain of 3000 B. An 80 per cent chance of gaining 4000, otherwise nothing

Problem 2
Which would you prefer:
A. A 0.1 per cent chance of gaining 5000 B. A certain gain of 5

Problem 3
Which of the following situations would you prefer:

Situation A:	Situation B:
100 million for certain	A 10 per cent chance of 500 million
	An 89 per cent chance of 100 million
	A 1 per cent chance of nothing

Problem 4
Which of the following situations would you prefer:

Situation C:	Situation D:
An 11 per cent chance of 100 million	A 10 per cent chance of 500 million
An 89 per cent chance of nothing	A 90 per cent chance of nothing

Problem 5
In addition to whatever you own, you have been given 1000. You are now asked to choose between:
A. A 50 per cent chance of 1000, otherwise nothing. B. 500 for certain

Problem 6
In addition to whatever you own, you have been given 2000. You are now asked to choose between:
C. A 50 per cent chance of losing 1000, otherwise nothing. D. Losing 500 for certain.

INTRODUCTION

Each problem shown above is a *risky* decision. Technically speaking, this means that particular outcomes occur only with a stated probability. Most economic and psychological accounts of risky decision making have taken a quantitative approach, and assume that when people make risky decisions they are trying to maximise something such as expected value or expected utility. However, economic theories fail to capture much of human decision-making behaviour. By contrast, prospect theory is a modification of expected utility theory that accounts for many of the violations of that theory.

Following on from prospect theory, I take a short excursion into the neuroscience of valuation to show how brain-imaging research is producing findings that relate to cognitive theories about decision making, including prospect theory.

In most everyday decisions probabilities are unknown. I describe two approaches to decision making under uncertainty. One of these is a two-stage model that combines support theory with a later development of prospect theory, known as cumulative prospect theory. The second approach is referred to as risk sensitivity theory, and has developed from the literature on optimal foraging in animals.

Many other theories of decision making have been proposed, some of them in opposition to prospect theory and some as a supplement to it. Process models of decision making are concerned with the stages of thinking that lead up to a decision. Two examples I describe here are decision field theory and the priority heuristic.

DECISIONS UNDER RISK (1): EXPECTED VALUE THEORY AND EXPECTED UTILITY THEORY

Value, utility, and the St Petersburg paradox

The start of this chapter showed a number of decision problems involving described probabilities and monetary outcomes. According to *expected value theory* the optimum action on such problems can be determined calculating the value of each possible outcome and weighting those outcomes by their probability of occurrence (i.e. multiply the outcomes by their probabilities). Then you choose the course of action with the highest *expected value*. For example, which of the following gambles would you prefer?

1. A 70 per cent chance of £100, otherwise nothing.
2. A 35 per cent chance of £250, otherwise nothing.

According to the criterion of expected monetary value, Gamble 2 is the better bet because it has an expected value of £87.50 (.35 × 250) whereas Gamble 1 has an expected value of £70 (.7 × 100).

However, Daniel Bernoulli (1954[1738]) observed that people do not behave as though they are maximising expected value. He described a game known as the St Petersburg paradox. Suppose I toss a fair coin and keep on tossing it until it lands on heads. If it lands on heads on the first toss, the game ends and you win one euro. If it does not land on heads until the second toss, then the game ends on the second toss and you win two euros. If it lands on heads on the third toss then the game ends and you win four euros. With every toss that the game continues the amount that I must pay you doubles. The question is, how much would you be willing to pay me in order to participate in this game?

Based on the concept of expected value you should be willing to pay me quite a large sum of money because the game itself has *infinite* expected value.[1] Bernoulli argued that a point would come in the game where further tosses of the coin would barely add to the *utility* that would accrue if the game ended at that point (where utility refers to pleasure or usefulness).

This idea – referred to as diminishing marginal utility for gains – is shown graphically in Figure 7.1. As a person's wealth increases, each extra unit of money adds utility but by less than the previous unit. This means that more wealth is always better, but by the same token an increase in wealth that would make a poor person very happy will not have the same utility for someone who is very wealthy. The property of diminishing marginal utility is consistent with Weber's Law, a psychological law that is normally applied to judgments of physical magnitude (see Box 7.1).

Bernoulli's theory explains why even bets are unattractive. Suppose two men are playing a game of chance in which there is

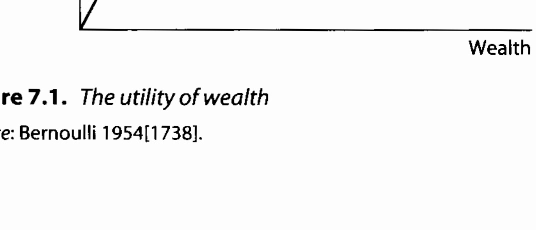

Figure 7.1. *The utility of wealth*
Source: Bernoulli 1954[1738].

BOX 7.1. WEBER'S LAW AND THE UTILITY FUNCTION

One intriguing aspect of Figure 7.1 is that this curve resembles the curves that are produced in magnitude estimation studies. For example, if we are asked to judge the heaviness of various weights, and our estimates are plotted against the actual weights, then we get a graph rather like Figure 7.1. That is, as the weights become heavier we are less able to accurately assess further increases. If, instead of plotting the raw amounts, we plot the logarithm of the magnitude estimate against the logarithm of the stimulus intensity we get a straight line. The relationship between magnitude estimate and stimulus intensity is that of a *power law*, as described by the equation $P = KS^n$, whereby perceived magnitude, P, equals a constant , K, times the stimulus intensity, S, raised to a power, n.

Some studies have asked people to make direct comparisons between two stimuli. Suppose George is asked to hold a weight in each hand. If he can tell the difference between a weight of 100 grams and a weight of 102 grams, but not a smaller difference, then his *difference threshold* for a 100 gram weight is 2 grams. However, George is unable to tell the difference between a 200 gram weight and a 202 gram weight, which is also a difference of 2 grams. In fact, his difference threshold for a 200 gram weight turns out to be 4 grams.

For each comparison, if we divide the difference threshold by the stimulus weight, we get the same figure:

$$2/100 = 0.02 \quad \text{and} \quad 4/200 = 0.02$$

The relationship between difference threshold and stimulus magnitude is known as *Weber's Law* and the constant that results is known as *Weber's Fraction* (after Ernst Weber, 1795–1878; see Goldstein, 2007). The observation that the utility function in Figure 7.1 resembles judgments about other stimuli suggests that people respond to money in the same way that they seem to respond to sensory stimuli.

a 50/50 chance of winning or losing. Each man is worth £10,000 and each bets £5000 on the game. For each man there is a 50/50 chance of ending up worth either £15,000 or £5000, with the expected outcome being (£15,000 + £5000)/2 = £10,000. Thus, the expectation is one of no change. However, the utility curve says that the loss of £5000 has greater negative utility than the gain of £5000 has positive utility – that is, the potential loss looms larger than the potential gain. Therefore, most people would be unwilling to take the bet – they are risk averse.

Whereas expected value theory says that rational decision makers should weight monetary outcomes by the probability of their occurrence, expected utility theory says that rational decision makers should weight the utilities of outcomes by their probability of occurrence. The concept of utility has the advantage of being applicable to things other than money, although it has the disadvantage of not being directly measurable. Nonetheless, techniques do exist for eliciting people's utility functions (e.g. Goodwin & Wright, 2004). The notion that people *are* utility maximisers has become widespread in economics and in evolutionary biology.

Three failures of expected utility theory

Risk-seeking versus risk-averse behaviour Although the utility function in Figure 7.1 suggests risk aversion, utility theory allows that some people may have a different utility function. Nonetheless, they are still considered rational as long as they attempt to maximise expected utility. In practice, though, people are widely assumed to be risk averse. However, contrary to this assumption, research reported by Kahneman and Tversky (1979, 1992) has indicated a fourfold pattern of risk attitudes depending on whether outcomes are gains or losses, and whether probabilities are small or medium-to-large.

Problems 1–6 below, first mentioned at the start of the chapter, are taken from Kahneman and Tversky (1979). The numbers in round brackets read as (outcome, probability) and the numbers in square brackets are the percentage of people choosing each option. In Problem 1, most people prefer a sure gain of 3000 over an 80 per cent chance to gain 4000. The expected value of the risky option is 3200, which is more than for the sure thing, hence people are risk averse in their choice. However, when the outcomes are losses rather than gains people mostly choose the risky option rather than the sure loss (Problem 1′). In this case, the risky option has the worse expected value (−3200) hence the majority choice is risk-seeking.

These two problems involve large probabilities. However, when the probabilities involved are very small (Problems 2 and 2′), then people are risk-seeking for gains and risk averse for losses.

Problem 1

(4000, .80) < (3000)
$N = 95$ [20] [80]

Problem 1′

(−4000, .80) > (−3000)
$N = 95$ [92] [8]

Problem 2

(5000, .001) > (5)
$N = 72$ [72] [28]

Problem 2′

(−5000, .001) < (−5)
$N = 72$ [17] [83]

Kahneman and Tversky (1992) confirmed this pattern of responses in a series of two-outcome gambles where there were no sure things (each outcome being probabilistic). Thus, the assumption of risk aversion is not supported.

The Allais paradox: evidence that people do not maximise expected utility Another problem with expected utility theory is that people do not always choose decision options that maximise expected utility. This was demonstrated by Allais (1953, 1990). Consider again Problems 3 and 4 from the start of the chapter:

Problem 3
Which of the following situations would you prefer:

Situation A:	Situation B:
100 million for certain	A 10 per cent chance of 500 million An 89 per cent chance of 100 million A 1 per cent chance of nothing

Problem 4
Which of the following situations would you prefer:

Situation C:	Situation D:
An 11 per cent chance of 100 million An 89 per cent chance of nothing	A 10 per cent chance of 500 million A 90 per cent chance of nothing

Most people say that they prefer A to B in Problem 3 and D to C in Problem 4 (Slovic & Tversky, 1974). However, this is not a choice that would be made by someone maximising utility. Why is this? Firstly, bearing in mind that there is no utility for a zero outcome we can express the preference for A over B as follows:

$$u(100 \text{ million}) > 0.89\ u(100 \text{ million}) + 0.1\ u(500 \text{ million})$$

Subtracting 0.89 u(100million) from both sides of the inequality leaves us with:

$$.11\ u(100 \text{ million}) > 0.1\ u(500 \text{ million})$$

However, if we describe the preference for D over C in the same fashion, we now get the reverse inequality:

$$.11\ u(100 \text{ million}) < 0.1\ u(500 \text{ million})$$

In fact, what I have just shown is that options C and D in Problem 2 are actually obtained by subtracting an 89 per cent chance of 100 million from both A and B. The fact that preferences change when a common element is removed from A and B violates an axiom of expected utility theory, known as Savage's (1954) sure-thing principle (see Box 7.2). This says that people's choices should be based upon those attributes of options that *differ*: attributes that are the same in both options should not affect choice. The

BOX 7.2. THE ALLAIS PARADOX AND SAVAGE'S SURE-THING PRINCIPLE

Another way of thinking about the Allais problem is to imagine the situations as lotteries in which the probabilities are represented by numbered balls (Table 7.1). A person taking part in one of these four lotteries would simply be choosing a ball at random from an urn. The interesting column in Table 7.1 is the last one, showing the outcome that would occur if a person were to select a ball numbered between 12 and 100 inclusive (representing 89 per cent probability). In Lottery A, a player would win 100 million if they picked a ball from 12 to 100, or any other ball besides. Similarly, in Lottery B a player will also win 100 million if they pick any ball from 12 to 100. Thus, when deciding whether to play Lottery A or B a person should not consider what might happen if they pick a ball from 12 to 100 because the outcome would be the same in each case (the sure-thing principle).

Table 7.1. *The Allais paradox represented as a lottery*

	Ball numbers		
	1	***2–11***	***12–100***
Situation 1			
Lottery A	100 million	100 million	100 million
Lottery B	Nothing	500 million	100 million
Situation 2			
Lottery C	100 million	100 million	Nothing
Lottery D	Nothing	500 million	Nothing

Now consider Lotteries C and D. Once again, for balls 12 to 100 the outcomes are identical for each lottery, except that now the outcome is to win nothing. Again, balls 12 to 100 do not need to be considered when deciding which lottery to play. However, if you now blank out this last column in Table 7.1 (equivalent to applying the sure-thing principle) you will see that the choice between A and B is exactly the same as the choice between C and D.

Rather than conclude that people are irrational, Allais himself seems to have regarded the 'anomalous' preferences as so compelling that he drew the conclusion that it is the theory, rather than people, that is at fault. Student participants do not find Savage's axiom especially compelling. A study reported by Slovic and Tversky (1974) found that most people continued to make paradoxical choices even after they had heard arguments advocating Savage's position. In a second study people heard arguments both for and against the sure-thing principle *prior* to making a choice. More people rated the Allais argument as more compelling than Savage's (51 per cent vs. 42 per cent). Curiously, given the preference for Allais's argument, more people made choices consistent with the sure-thing principle (61 per cent) rather than inconsistent with it (35 per cent). However, comments by some subjects indicated that their choices satisfied the sure-thing principle because they felt Allais's recommendation was too conservative, and not because they found Savage's argument more compelling.

preference pattern in Allais's problem also indicates, of course, that people are not maximising expected utility.

Kahneman and Tversky (1979) described performance on the Allais problem as representing a *certainty effect*. In other words, a switch from certainty to uncertainty (or vice versa) exerts a particularly large effect on people's preferences.

Evidence that people do not integrate prospects with existing assets A third way in which expected utility fails to describe how people think about decisions is its assumption that they integrate possible outcomes (prospects) with their current assets. Kahneman and Tversky presented two different groups of participants with one of the following problems:

Problem 5
In addition to whatever you own, you have been given 1000.
You are now asked to choose between

A.	(1000, .50),	and	B. (500)
$N = 70$	[16]		[84]

Problem 6
In addition to whatever you own, you have been given 2,000.
You are now asked to choose between:

C.	(−1000, .50),	and	D. (−500)
$N = 68$	[69]		[31]

The choice percentages shown in square brackets show a reflection effect. However, a person behaving according to utility theory should either choose A and C or B and D. Utility theory says that we should integrate decision outcomes with our current assets. If we do this we see that the outcomes for both A and C are (2000, .50; 1000, .50) and the outcomes for B and D are (1500). It seems that people are not integrating their assets and outcomes in this way, a phenomenon that Kahneman and Tversky refer to as the *isolation effect*.

DECISIONS UNDER RISK (2): PROSPECT THEORY

The value function and the probability weighting function

Prospect theory can be considered as a psychological variant of subjective expected utility theory. The process of considering a decision begins with an *editing stage*, in which decision makers structure the decision problem in such a way as to simplify subsequent evaluation and choice (the *evaluation stage*). One important way in which they do this is by coding potential outcomes as gains and losses relative to some reference point. This reference point is often the status quo, but may also be an expectation or an aspiration – an important difference from utility theory, in which potential outcomes are always evaluated in relation to the current

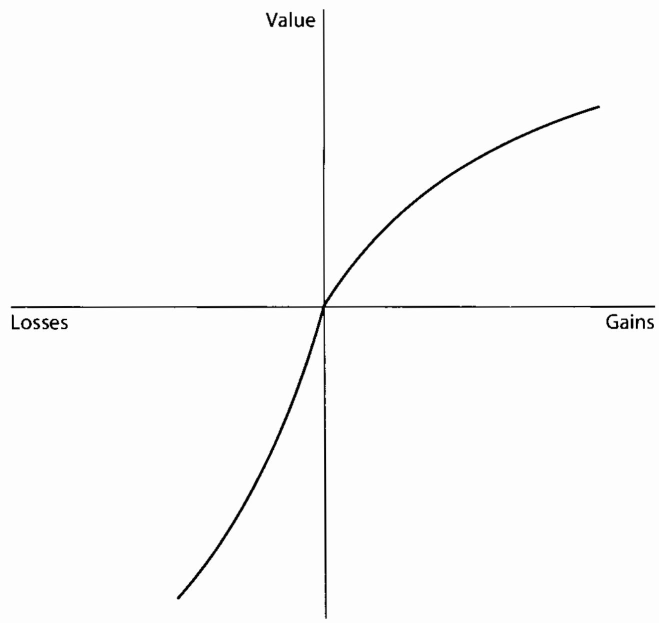

Figure 7.2. *A hypothetical value function in prospect theory*

asset position. The demonstration of the isolation effect shows that people do not always integrate potential outcomes with existing assets, as required by utility theory. The nature of the reference point – and hence, whether outcomes are evaluated as gains or losses – can also be affected by the description of the problem (see the section below on *framing effects*).

At the *evaluation phase*, both a value function and a weighting function are applied to prospects. The S-shaped curve in Figure 7.2 is a representation of the value function (value is pretty much the same as utility, except for its being assessed relative to a reference point). This curve is said to be concave for gains and convex for losses, capturing the notion that people are more sensitive to changes occurring near the reference point than to those further from it.

In discussing the utility function in expected utility theory, I noted that losses loom larger than gains (the reason that people usually turn down even bets). In prospect theory, the cause of this phenomenon is reflected in a steeper curvature of the value curve for losses compared to gains. However, as we have already seen, in prospect theory possible changes in fortune are not necessarily evaluated in relation to current assets. Thus, a potential objective gain of a certain amount could still be evaluated as a potential loss if it falls short of an aspiration or expectation. Furthermore, according to prospect theory initial losses are felt more keenly than subsequent losses of the same amount – this does not follow from expected utility theory.

Prospect theory also adopts a decision-weighting function, π (Figure 7.3). One reason that this is needed is to account for the large effect on preferences when outcomes switch from certain to uncertain (the *certainty effect*), as in the Allais paradox. On that particular problem, the switch from certainty to uncertainty led to a preference reversal. The weighting function is also needed to account for the big impact that small probabilities have on preferences (perhaps accounting for the popularity of lottery tickets and insurance).

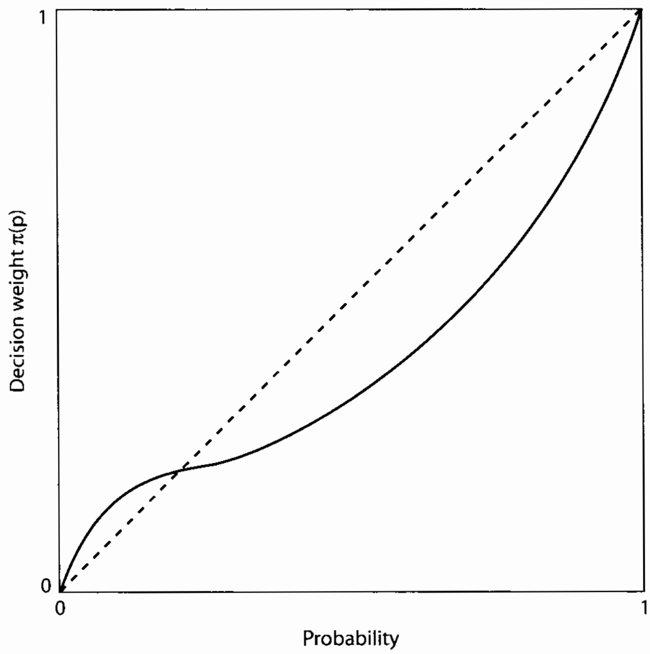

Figure 7.3. *The decision weighting function in prospect theory*

The weighting function has a reverse S-shape[2] that shows how stated probabilities are transformed, such that decisions are highly affected by small probabilities, whereas they are much less sensitive to medium-to-high probabilities (except at the boundary of uncertainty and certainty). There is evidence that the curvature of the weighting function is more pronounced when outcomes arouse affective feelings. For example, Rottenstreich and Hsee (2001) found that people attached greater value to shifts from certainty to a 99 per cent probability when outcomes were (for example) electric shocks rather than cash penalties. By contrast, people attached little value to the difference between a 99 per cent chance of an electric shock and a 1 per cent chance, whereas they attached considerably more value when the difference involved the chance of a financial penalty.

In general, weighted probabilities sum to less than 1, a property that is referred to as *subcertainty*. It should be noted that the decision-weighting function is not a measure of what people believe the probabilities to be. After all, in the problems shown above the probabilities are actually stated for the participants. Decision weights simply measure the *impact* that probabilities have on decisions.

Prospect theory not only explains a wide range of phenomena observed in laboratory experiments, but has also been applied to a range of real-world settings (Box 7.3). In this sense, it is the most

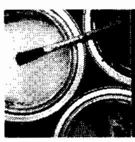

BOX 7.3. APPLICATIONS OF PROSPECT THEORY

Applications of prospect theory have been investigated in many domains, both in the laboratory and in real-world settings. These will be explored in more depth in the next chapter, but a number of illustrative phenomena are described below.

As we saw earlier, the overweighting of low probabilities helps explain the popularity of both lotteries and insurance. Similarly, it turns out that there is a bias towards betting on long-shots in horse races; that is, the percentage of time that long-shots win is less than the percentage of money bet on them (Camerer, 2000). There is also an end-of-the-day effect in horse racing, whereby betters tend to shift their bets away from favourites and towards longshots in the last race of the day (Ali, 1977; McGlothlin, 1956). Expected utility theory cannot easily explain this shift, because a bettor who integrated his wealth would not treat the last race of the day as any different from the first race of the next outing. However, prospect theory assumes that people have a minimal target (or reference point) of breaking even at the end of the day. Anything below this target is regarded as a loss and triggers risk-seeking behaviour in an attempt to break even. Diminishing sensitivity to successive losses means that the cost of the last bet is relatively trivial compared to what has already been spent.

A similar targeting effect has also been observed in New York Cab drivers (Camerer *et al.*, 2000). According to standard labour-supply theory, drivers should maximise their profits by working longer on good days and quitting earlier on bad days. In fact, for drivers with several years' experience there was no correlation between daily wage and hours worked, whereas inexperienced drivers quit earlier on good days. Thus, inexperienced drivers fail to maximise their earnings because they quit once they have met their daily target. Over time, these drivers either quit or shift towards driving about the same number of hours each day.

Daily targeting of this sort is computationally simpler than taking the longer perspective. It also helps mitigate certain self-control problems. For example, a driver who takes a longer perspective on his earnings might be tempted to quit early today and make up for the shortfall tomorrow. However, when tomorrow arrives he might be tempted to do the same again, thus getting into an endless cycle of putting off the hard work until some future time.

Several authors (reviewed in Camerer, 2000) have observed a *disposition effect* in investment. This refers to a tendency of investors to hold on for too long to stocks that have lost value, but to sell stocks that have risen in value. The disposition effect has been demonstrated experimentally and in real-life settings, and appears to be a consequence of loss aversion: investors are more sensitive to their losses than to their gains, and are more willing to gamble to recoup those losses. Of course, it could be that investors hold on to failing stocks because they expect them to 'bounce back'. However, Odean (1998) reported that one year after winning stocks had been sold, they were still out-performing the losers that had been retained.

successful account of human decision making so far proposed. Nonetheless, other research has identified some aspects of decision performance that appear to conflict with prospect theory (e.g. Birnbaum, 2006; see also the section on *risk sensitivity theory* below). Some researchers have also criticised prospect theory for its lack of detail about the underlying cognitive processes (see the final section below on *process theories of decision making*).

Framing effects in decision making

According to the invariance axiom of utility theory *The outcomes and associated probabilities are all that is required to determine a decision maker's preference between uncertain events.* In other words, irrelevant changes of wording in the description of a decision should not affect our preferences regarding that decision. We have already seen that people violate this principle, however (Problems 5 and 6 above). By coding the same outcomes as either a gain or a loss, it is therefore possible to change the way that people evaluate those outcomes.

In another demonstration of framing, Kahneman and Tversky (1984) asked their participants to imagine that the USA was preparing for the outbreak of an unusual Asian disease, expected to kill 600 people. Two alternative programmes for combating the disease had been proposed, for which the exact scientific estimate of the consequences were as follows:

If programme A is adopted, 200 people will be saved (72 per cent)

If programme B is adopted, there is a one third probability that 600 people will be saved and a two thirds probability that no people will be saved (28 per cent)

The numbers in parentheses represent the proportion of respondents choosing each option and indicate that most people preferred option A, which could be characterised as a riskless option as it guarantees that 200 people will be saved. In comparison, the possibility of loss for option B makes the gamble unattractive.

However, other participants were presented with the same scenario and asked to choose between option C and option D:

If programme C is adopted, 400 people will die (22 per cent)

If programme D is adopted, there is a one third probability that nobody will die and a two thirds probability that 600 people will die. (78 per cent)

Of course, programmes C and D are actually the same as programmes A and B except that now the outcomes are framed in terms of the numbers of lives that might be lost. We can see that now the risky option (D) is more popular than the riskless option.

One early study reported that experts were also susceptible to framing effects. McNeil *et al.* (1982) found that both physicians and laypeople altered the choice of surgery or radiation therapy as a treatment for lung cancer depending on whether likely outcomes were described in terms of survival rates or mortality rates. Nonetheless, in actual medical situations this type of framing effect appears not to be robust and other factors appear to be more important (Christensen *et al.*, 1995; Siminoff & Fetting, 1989; for a review of medical decision making see Hamm, 2003). However, in the financial domain framing effects appear to be much more robust. Richard Thaler (1980) noted that lobbyists for the US credit card industry had an interesting response when it appeared that a bill would be passed to allow stores to charge higher prices to credit card users. They preferred that this price difference be described as a cash discount rather than a credit card surcharge. Similar findings will be discussed further in the next chapter.

There is some evidence that positive and negative framing are associated with different levels of cognitive processing. Dunegan (1993, Experiment 1) presented a project-funding scenario to 128 members of an international company engaged in the development of high-technology engineering systems. The scenario tested the extent to which participants would be willing to jeopardise their ability to fund certain opportunities by allocating further funding to an existing project that was already behind time and over budget, but which could most likely still be successfully completed. The positive frame stated that *Of the projects undertaken by this team, 30 of the last 50 have been successful*, and the negative frame stated that *Of the projects undertaken by this team, 20 of the last 50 have been unsuccessful.*

Dunegan found that the negative framing group allocated less money on average to the current project and they rated themselves as less likely to actually fund such a request. In addition, the negative framing group perceived greater risk, were more disappointed in the project, and were more concerned about minimising losses. These measures, together with some others, accounted for a significant proportion (45 per cent) of the variance in funding allocations among the negative framing group, but were not predictive of funding allocations in the positive framing group. A follow-up study, using students, found that people in the negative framing condition had a greater disparity between their images of the current state and their images of the goal state. Dunegan argued that the negative framing condition induced a negative affective state, leading to a more controlled mode of cognitive processing.

Consistent with the idea that negative frames lead to more cognitive processing, other studies have found that negative frames are associated with a slower response time (Payne *et al.*, 1993; Gonzalez *et al.*, 2005). In the study by Gonzalez and colleagues, student volunteers responded to a series of framing problems while inside a functional magnetic resonance imaging (fMRI) scanner. As usual, the predominant choices were the sure-thing option in the positive frame and the risky option in the negative frame. For positively framed problems less brain activity was evident when the sure-thing option was chosen than when the risky option was chosen. For negatively framed problems there was no difference in brain activity for the sure-thing and risky options.

Gonzalez *et al.* found that risky choices were associated with higher levels of activity in brain regions associated with imagery and working memory function (specifically, the frontal and parietal lobes).[3] They suggested that there is an interplay between affect and cognition in framing problems. That is, people generally try to minimise cognitive processing but when a problem induces negative feelings then the level of cognitive processing increases. Thus, in the positive frame the sure-thing option tends not to arouse negative feelings and is relatively quickly accepted. However, in the negative frame both options arouse negative feelings and lead to more processing.

Another fMRI study of framing has also indicated a role for affective influences (De Martino *et al.*, 2006). This study endowed participants with an amount of money and then asked them to choose between a sure thing (gain or loss, where the loss was less than the endowment, hence an overall gain was involved) or a risky option involving a chance to keep all or lose all. Choice of the sure thing in the gain frame and the gamble option in the loss frame were both associated with greater amygdala activity. This is an area of the brain associated with affective processing. However, when people made choices that ran counter to their general tendency – typically, choosing the gamble in the gain frame and the sure thing in the loss frame – there was increased activity in the anterior cingulate cortex (ACC), an area believed to be responsible for conflict monitoring and cognitive control. Taken together, these results indicate opposing influences between an emotional amygdala-based system and an analytic ACC system.

Indications of dual-processing mechanisms are consistent with the evidence that framing effects are moderated by cognitive ability and cognitive motivation. For example, in one study smokers at a public event read smoking-cessation messages that were either framed in terms of the lives that could be saved if smokers quit the habit or in terms of the numbers that would die if smokers continued to smoke. They were then asked about their own intentions to quit. People who enjoyed engaging in effortful thinking were unaffected by the framing manipulation, but less motivated thinkers registered a greater intention to quit after reading the gain-framed message (Steward *et al.*, 2003). In another study, Stanovich and West (1998) found that students of higher cognitive ability – as indicated by their SAT scores – were less susceptible to framing on the Asian disease problem described above.

Rationales for the value and probability weighting functions

Although the value function and probability weighting functions account for a great deal of decision behaviour, it is not clear why these functions are the way they are. However, the concept of declining marginal utility (see Figure 7.1 and the top right-hand section of Figure 7.2) has a parallel in nature. For example, Harder and Real (1987; see also Real, 1996) showed that the rate of net energy gain to bees shows diminishing returns as the nectar reward size increases. Plotted graphically, this is of the same form as Figure 7.1.

Evidence that people may be sensitive to the structure of gains and losses in the environment has been provided in a theory of *decision by sampling* (Stewart *et al.*, 2006). This theory was inspired by other approaches that involve memory sampling, such as the MINERVA-DM theory that we encountered in Chapter 3. According to this theory, when we are faced with a decision problem involving monetary amounts and probabilities (for example), we compare these attribute values with previously encountered values sampled from memory. The contents of memory are, in turn, assumed to reflect the structure of the world.

Stewart *et al.* investigated the frequency with which differing monetary amounts were manually entered into or withdrawn from people's bank accounts. Both credits and debits were found to follow a power law. There were many small gains and losses, and relatively fewer larger gains and losses. Stewart *et al.* argue that the preponderance of small gains and losses means that these are more frequently sampled from memory, hence people are more sensitive to smaller outcomes. This is consistent with the curvature of the value function in prospect theory. There were also relatively more small losses compared to small gains. This asymmetry is consistent with the observation that losses loom larger than gains.

Other evidence points towards probabilities being compared with attribute values being retrieved from memory. For instance, larger working memory spans correlate with less subadditivity in probability judgments (Dougherty and Hunter, 2003a, 2003b). Stewart *et al.* also review several sources of evidence indicating an S-shaped function for probabilities whereby both smaller and larger probabilities are over-represented in the environment relative to middle-range probabilities. For example, Gonzalez and Wu (1999) conducted a study to investigate the shape of the probability weighting function. When the relative ranks of these probabilities were plotted against the probabilities themselves the result was an S-shaped curve, indicating that small probabilities are overestimated and large probabilities underestimated.

THE NEUROSCIENCE OF VALUATION

Montague (2006) has made the point that mobile organisms need to have an internal model of the world, and that this requires assigning values in order that goals can be chosen and prioritised, and so feedback can aid learning. Furthermore, it is computationally more efficient to have a single mechanism for valuing different types of goals and stimuli:

> Without internal currencies in the nervous system, a creature would be unable to assess the relative value of different events like drinking water, smelling food, scanning for predators, sitting quietly in the sun, and so forth. To decide on an appropriate behaviour, the nervous system must estimate the value of each of these potential actions, convert it to a common scale, and use this scale to determine a course of action. This idea of a common scale can also be used to value both predictors and rewards. (Montague & Berns, 2002, p.276)

Studies of primates have shown that midbrain dopaminergic neurons[4] play a crucial role in the valuation of rewards and in developing expectancies (e.g. Romo & Schultz, 1990). For example, when a monkey receives an unexpected squirt of juice there is a sudden activity 'spike' in dopamine neurons in the ventral tegmental area. If, however, the monkey receives a series of trials on which the juice squirt is preceded by a flash of light one second earlier, then the dopamine spike eventually stops appearing after the reward but instead appears after the light flash. Once the light–juice relationship has been learnt, any occasional non-appearance of the squirt following the light results in a fall in dopamine activity below the normal baseline.

With human participants, the dopamine system responds not just to actual rewards but to symbolic information. For instance, young men show a greater dopamine response to pictures of sports cars, which have a high status, than to other types of cars (Erk *et al.*, 2002). In another series of studies (McClure *et al.*, 2004), a cup of Coke led to greater brain activity when the cup was labelled 'Coke', and to greater liking, than when the cup was unlabelled (even though participants could not be absolutely certain that they were really drinking Coke). This brand name effect did not work for Pepsi: brain activity and liking were not increased when a Pepsi drink had a Pepsi label attached to it (even though people did not have a preference between Pepsi and Coke in a blind taste test).

These kinds of studies indicate that a single mechanism is used for evaluating different kinds of stimuli. Furthermore, both people and non-human primates develop expectations based on experience and evaluate outcomes in relation to those expectations. In the case of people purely symbolic information can also be rewarding.

Other studies have examined people's responses to gambles. Breiter *et al.* (2001) used fMRI to examine activity in brain regions into which dopaminergic neurons project – namely the orbito-frontal cortex, the nucleus accumbens, the amygdala, the sub-lenticular extended amygdala (SLEA) of the basal forebrain, and the hypothalamus – and which project back to the area from which these projections arise. During a 'prospect' phase participants saw an arrow spin around a disk divided into three regions representing different monetary outcomes. When the arrow stopped on one of the three regions, that area began to flash, marking the start of the 'outcome' phase. Three spinners were used, each of which had zero as one of its outcomes. A good spinner had two positive outcomes in addition to the zero outcome, a bad spinner had two negative outcomes, and an intermediate spinner had a positive and a negative outcome.

During the prospect phase, the SLEA and the orbital gyrus appeared to track expected value: they both showed the greatest activity in response to the good spinner and the least activity in response to the bad spinner. Responses to actual outcomes increased monotonically with value in the nucleus accumbens, SLEA, and hypothalamus.

Another aspect of decision behaviour that has been investigated with fMRI is loss aversion. Tom *et al.* (2007) reported evidence that this phenomenon is not the result of negative affect. When participants were asked to accept or reject gambles involving a 50/50 chance of gaining one amount or losing another, brain areas associated with negative emotions showed no increase in activity as the size of potential losses increased. Rather, the processing of gains and losses occurred mainly in the same brain areas (the ventral striatum and ventromedial prefrontal cortex), indicating an aggregate representation of decision utility. Activity in these areas decreased as potential losses increased. Loss aversion in behaviour was strongly correlated with neural loss aversion.

However, Shiv *et al.* (2005) *did* find evidence for the role of emotions in *myopic* loss aversion. In myopic loss aversion people who are faced with a series of bets do not think about their long-run potential gains and losses, but instead evaluate bets one at a time, thus continually rejecting advantageous gambles due to the possibility of a loss. The Shiv *et al.* study was slightly different from that of Tom *et al.*. Whereas Tom *et al.* did not resolve the gambles during the process of scanning, Shiv *et al.* did do so. Furthermore, each gamble was the same. Specifically, participants played 20 rounds in which the choice was always whether to invest $1 in a gamble offering a 50 per cent chance of gaining $2.50 and a 50 per cent chance of gaining nothing. The optimal strategy is, of course, to invest in each gamble. The participants were people with damage to brain areas involved in processing emotion (the target patients), patients with brain damage in non-emotion areas, and normal controls. The most optimal players were the target patients. They invested 84 per cent of the time, compared to 61 per cent for the brain damage controls, and 58 per cent for the normal controls. All three groups of participants tended to invest on the earliest rounds, but for both groups of control patients this declined across the course of the rounds. This decline occurred regardless of the experience of winning or losing rounds, although losing appeared to make people even less likely to invest.

DECISION MAKING UNDER UNCERTAINTY

As we have seen, prospect theory posits a weighting function that transforms stated probabilities. However, for most events that we encounter in real life there are no stated probabilities. In this section I shall describe two approaches to uncertain decisions, one based upon an extension of prospect theory and another based upon a model of decision making from the literature on behavioural ecology.

Support theory and cumulative prospect theory: a two-stage model of decision making

To explain uncertain decision making, Kahneman and Tversky (1992) developed a modified version of prospect theory, known as cumulative prospect theory. Later, this was combined with support theory (Tversky & Koehler, 1994; see Chapter 3) into a two-stage model (Fox & Tversky, 1998; see also Fox & See, 2003).

In the two-stage model, people first form an assessment of the probabilities of uncertain events. Secondly, they apply a weighting to these probabilities, adjusted for their own perceived knowledge of the domain in question. These weights are then combined with the output of prospect theory's value function. Consistent with support theory, Fox and Tversky (1998) found that participants priced prospects higher when they were given a more specific description. For example, participants judged a 78 per cent likelihood that *a team* from the Eastern Conference would win the National Basketball Association's 1996 playoffs. However, a separate group of participants assigned a 90 per cent probability that *one of the four leading teams* from the Eastern Conference would win the playoffs. This ordering of likelihoods is the reverse of what should rationally be the case: the four leading teams are merely a subset of the Eastern Conference so it is less likely that the winner of the playoffs will be among them than among the entire Eastern Conference.

Other participants were asked to state their certainty equivalent for a prospect that offered $75 if one of these events came about.[5] The median certainty equivalent was lower for the prospect of an Eastern Conference team winning ($50) than for one of its four leading teams winning ($60). These certainty equivalents are indicative of the weightings that people attach to these events and, as with the group who assessed likelihoods, they are the opposite of what should be rationally expected.

The weightings that people attach to judged probabilities are also modified to take account of people's perception of their knowledge about the domain in question. This accounts for the phenomenon of *ignorance aversion*. Heath and Tversky (1991) found that people preferred to bet on vague events that they felt knowledgeable about rather than chance events that they considered equally probable. For example, people who felt knowledgeable about football preferred to bet on uncertain football events rather than chance events. However, when people lacked knowledge of football they preferred to bet on chance events.

Figure 7.4 is a visual representation of Fox and Tversky's (1998) two-stage model of uncertain decision making. This incorporates both the value function and the decision-weighting function from prospect theory, except that decision weighting is now influenced by judged probabilities and modified by the decision maker's attitude towards the source of uncertainty.

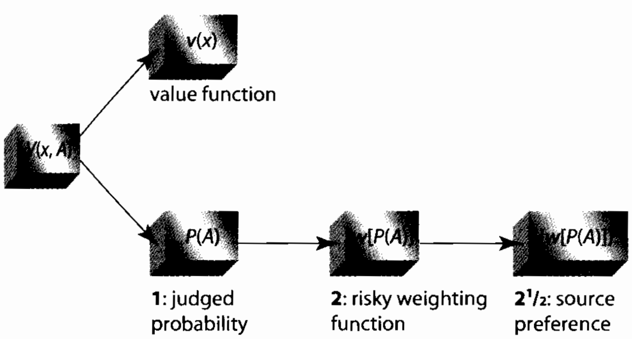

Figure 7.4. *Visual depiction of the extended two-stage model. $V(x, A) = v(x)\,(w[P(A)])^{\theta}$ where $V(x, A)$ is the value of the prospect that pays \$x if event A obtains (and nothing otherwise), v(.) is the value function for monetary gains, P(.) is judged probability, w(.) is the risky weighting function, and θ is the source preference parameter.*

Source: Fox & See, 2003.

Risk-sensitivity theory

Research into optimal foraging in animals has focused closely on animals' experience of their environments. This literature has given rise to a class of normative models known as risk sensitivity theories which, as we shall see shortly, have also been applied to humans. These theories 'construe risk-sensitivity as the response of organisms whose goal is the maximization of Darwinian fitness in stochastic environments' (Weber *et al.*, 2004, p.430).

Consider an animal that is foraging for food. It has the choice between two food sites. Both sites offer the same level of food on average, except that over time one site shows a lot more variation around this average than does the other site. What should the animal do? The *energy budget rule* (Caraco, 1980; Stephens, 1981) says that the animal's decision should depend on its current energy state. If the animal requires few resources, then it is better off going to the low-variation site, where it will most likely satisfy its requirements, before turning its attention to other matters (such as mate acquisition). If it goes to the high-variation site there is a greater chance that it will not be able to satisfy its food requirements. On the other hand, if the animal's food requirements exceed the average amount offered at the two sites, then it is better off following the risky strategy of choosing the high variation site. There, at least, a higher possibility exists of obtaining the necessary resources. Note that this prediction differs from the widespread assumption in the human research literature that people prefer certainty over ambiguity (Ellsberg, 1961; Slovic & Tversky, 1974).

The energy budget rule also predicts how organisms should behave where time delays are involved. It takes time to travel to a food source, so in choosing between two equivalent food locations an animal needs to take into account its current energy requirements and the nature of the delay in reaching each location. The energy budget rule predicts that an organism should prefer a fixed delay over a variable delay when the animal will not starve during the fixed delay period (holding food equivalent across delays). However, if the organism *is* likely to starve before the fixed delay expires, then it should opt for the variable delay because this offers the only possibility of staying alive long enough to obtain food.

Supporting the energy budget rule, Caraco found that wild-caught birds behaved in exactly this fashion when two seed sources were arranged in a laboratory setting (Caraco, 1981, 1983; Caraco *et al.*, 1980). The energy budget rule has also predicted foraging behaviour in several other species (Kacelnik & Bateson, 1996). However, not all studies have provided supportive evidence (Kacelnik & Bateson, 1996; Shafir *et al.*, 1999). One problem appears to be that it is difficult in practical terms to operationalise an animal's energy budget, because there are many things that may affect this such as ambient temperature, predation risk, and so on (see Pietras *et al.*, 2003, and also Soto *et al.*, 2005). Recognising these difficulties, two different approaches have developed, as described below.

One approach focuses on studying human participants in order that variables of interest are more tightly controlled. In one study human participants undertook several blocks of trials within which they had to try and earn points – to be translated into cash – within a limited time period (Pietras *et al.*, 2003; see also Pietras &

Hackenberg, 2001). They did this by choosing one of two flashing keylights to press. Pressing a keylight five times consecutively would turn off the other keylight, and after a delay the chosen keylight would stop flashing but stay illuminated, indicating that it had been chosen. One keylight had a fixed delay of 10 seconds, the other key had an equiprobable delay of 2 seconds or 18 seconds. A counter visible to the participants kept a record of the cumulative delay time. At the end of each block, if the cumulative delay did not exceed a certain threshold then the participant would be rewarded with 10 points (worth 25¢).

During positive budget conditions, the threshold was set at 50 seconds, so exclusive preference for the fixed option would always result in points earnings, whereas exclusive preference for the variable option would only result in earnings half the time. During negative budget conditions the threshold was set at 40 seconds or 32 seconds, so that exclusive preference for the fixed option would never result in points earnings but exclusive preference for the variable option would result in earnings with $p = 0.19$. Pietras *et al.* found that participants strongly preferred the fixed delay and variable delay, respectively, under these conditions. They also found that people were sensitive to the specific outcomes that they encountered; for example, under the negative budget people might switch to the fixed option if they had been fortunate enough to experience several short delays with the variable keylight.

Another study introduced the concept of *need* into simple gambles involving ambiguity (Rode *et al.*, 1999, Experiment 4). In this experiment participants aimed to blindly draw a threshold number of black balls from a box in order to gain a chance of winning a prize. Participants could choose to draw from a *risky* box with a specified number of black and white balls or from an *ambiguous* box with an unknown distribution of black and white balls. Within the conditions of the experiment the threshold number of black balls was varied, as was the known distribution of black and white balls.

Participants were considerably more likely to select the ambiguous box when the threshold exceeded the expected value of the risky box than when it did not. People were also sensitive to probabilities: they increasingly chose the ambiguous box as the probability of drawing black from the known box fell.

Other studies indicated that ambiguity aversion does not depend on the direct comparison of an ambiguous and a known option (Rode *et al.*, 1999, Experiments 1 and 2) and that people avoid options with high outcome variability regardless of whether probabilities are explicitly stated or not (Experiment 3). Overall, these results are consistent with optimal foraging theory.

A second approach to risk sensitivity has recognised that many studies have not taken into account individuals' *perceptions* of variability or risk. For instance, a price reduction of $100 sounds great when buying a $200 pen, but trivial when buying a $20,000 car (Weber *et al.*, 2004). This is consistent with our earlier discussion of the utility function and Weber's Law.[6] The *relative* variability of risky choice alternatives can be measured by dividing the standard deviation of outcomes by their mean (For those who haven't studied statistics: the standard deviation is simply a particular measure of variation. Don't worry if you're unfamiliar with this.) The resulting measure is referred to as the *coefficient of variation* (CV). As well as having greater psychological plausibility, the CV has the advantage of being dimensionless; that is, the unit of measurement is cancelled out, so it is possible to make comparisons across different domains (this approach is interesting in light of the evidence that humans and other animals use a single internal scale for value; see the earlier section on the neuroscience of valuation).

In a review of animal studies, Shafir (2000) found that the CV was a better predictor of risk sensitivity than was variance. Weber *et al.* (2000) reviewed human studies of decision making and found that the CV was a slightly stronger predictor of risk taking than was variance or standard deviation in both the gain domain and the loss domain. However, the results were not as strong as for the animal studies, possibly because the human studies did not call upon personal experience with the domain. To address this problem, Weber *et al.* examined choice behaviour in a card game where participants could choose cards from a constant payoff deck and a variable payoff deck. They found that risk aversion increased with the CV, but showed no relationship with the variance of outcomes. Expected value also predicted choice proportions, but CV predicted above and beyond expected value.

In summary, risk sensitivity theory offers an interesting perspective on decision making that differs from expected utility theory and prospect theory in that people are concerned with the variance of some outcome rather than some average value. However, there has still been relatively little research with human participants, although some of the early results described here are encouraging.

PROCESS MODELS OF DECISION MAKING

Some authors have tried to say more about decision making as a process. In *decision field theory* (Busemeyer & Johnson, 2004) the deliberations of a decision maker involve switching attention between alternatives and thinking about their consequences. In doing so, an overall feeling of desirability in relation to each action accumulates. Once a threshold is passed, then the successful action is chosen. Busemeyer and Johnson show how their theory provides an alternative account of many of the phenomena that prospect theory was designed to explain.

By contrast, Birnbaum (e.g. 1997, 2006) has identified several new paradoxes that are not accounted for by cumulative prospect theory, but which he explains in terms of a transfer of attention exchange (TAX) model.

In the rest of this section I shall describe in some detail another attentional account of decision making under risk. This account has been proposed by Brandstätter *et al.* (2006). These authors suggested that people apply a simple heuristic when choosing in a described decision task. They refer to this as the *priority heuristic*. This heuristic is based on the following steps:

- *Priority Rule.* Go through the reasons in the order: minimum gain, probability of minimum gain, maximum gain.

- *Stopping Rule.* Stop examination if the minimum gains differ by 1/10 (or more) of the maximum gain; otherwise, stop examination if probabilities differ by 1/10 (or more) of the probability scale.
- *Decision Rule.* Choose the gamble with the more attractive gain (probability).

In cases where the prospects are all negative, these rules are the same except that 'gains' are replaced by 'losses'.

To see how the priority heuristic works, let us revisit the Allais paradox (Problems 1 and 2 from earlier in the chapter).

Problem 1

Which of the following situations would you prefer:

Situation A:	Situation B:
100 million for certain	A 10 per cent chance of 500 million An 89 per cent chance of 100 million A 1 per cent chance of nothing

Problem 2

Which of the following situations would you prefer:

Situation C:	Situation D:
An 11 per cent chance of 100 million	A 10 per cent chance of 500 million
An 89 per cent chance of nothing	A 90 per cent chance of nothing

In Problem 1 the decision maker's aspiration level is 50 million; that is, one tenth of 500 million (the maximum gain). The minimum outcomes are zero (in B) and 100 million (in A). Because 100 million exceeds the aspiration level the sure thing in Situation A is chosen.

In Problem 2 the minimum gains are both zero (with probabilities of .89 and .90), thus they fall short of the aspiration level. The maximum gains are 100 million (in C) and 500 million (in D). Therefore, the priority heuristic predicts the choice of D.

The priority heuristic also predicts the reflection effect, as well as the observation of risk aversion for low probability losses and risk seeking for low probability gains. Furthermore, Brändstatter *et al.* showed that the priority heuristic outperformed prospect theory, as well as several other theories and heuristics, when used to predict choices on a series of gambles taken from the literature, as well as randomly chosen gambles. Finally, the authors conducted a study in which participants' response times were recorded on a series of choices of gambles. This found that response times were longer for problems where the priority heuristic predicted that more reasons would need to be examined.

Although these data are impressive, they must nonetheless be regarded with some caution. For example, the stopping rule seems fairly arbitrary, although the authors justify it on the basis of certain numbers being prominent within the decimal system. Indeed, animals have also been observed to violate the utility axioms on a version of Allais's problem (Battalio *et al.*, 1985; Kagel *et al.*, 1990), yet they are not – presumably – operating according to a decimal system. The response time data may also be interpretable in terms of a different process than a search for reasons. For instance, it may relate to a process involving decision weighting. There are also phenomena that the priority heuristic does not address, such as framing effects, the influence of affect on the weighting of probabilities, and some of the real-world phenomena that prospect theory successfully predicts.

SUMMARY

Faced with a choice between risky prospects, expected value theory states that each potential outcome should be weighted by the probability of its occurrence. The expected value of any given alternative is the sum of its weighted outcomes. The rational decision maker should choose the alternative with the highest expected value.

Bernoulli's St Petersburg paradox demonstrated a problem with the concept of expected value, leading him to develop the idea that what people are actually trying to maximise is expected utility. In the 20th century, a series of axioms was developed, which were said to underpin expected utility theory. A rational decision maker should follow these axioms if he wishes to make decisions consistent with his beliefs.

However, much psychological research has found many instances where people do not behave according to the axioms of expected utility theory. Prospect theory was developed as a theory of how people actually make decisions between two alternatives under risk. It recognises that the reference point against which gains and losses are defined may be an aspiration or expectation rather than the status quo. Furthermore, small probabilities tend to be overweighted in decision making, whereas medium-to-high probabilities are underweighted.

Key to any decision making is the concept of valuation. Recent evidence indicates that a single internal scale is used for valuing different types of outcomes, potential outcomes, and predictive cues. The brain's dopamine system appears to be the mechanism underlying the valuation process.

In order to explain decision making under uncertainty cumulative prospect theory was developed, and then combined with support theory into a two-stage model. People make an assessment of probability, which is then weighted and modified according to the decision maker's own knowledge of the domain in question. An alternative approach to uncertain decisions comes from optimal foraging theory. This specifies that choices are determined by a combination of the need of the decision maker, the mean expected outcome of each option, and the variance of each option's outcomes.

Finally, some approaches to decision making place greater emphasis on the stages of processing involved. One such approach is decision field theory, which emphasises switches in attention over time between different elements of the decision problem. Another approach is that of the priority heuristic, which argues that decisions under risk do not involve the computation of expectations; rather, people give sequential attention to outcomes and probabilities according to an order of prioritisation.

QUESTIONS

1. Do you think it makes sense for people to be more sensitive to losses than to gains? Why?
2. In what ways have people been observed to violate expected utility theory?
3. Compare and contrast prospect theory with expected utility theory.
4. What is the energy budget rule?
5. Construct alternative forms of a message framed, alternately, as a gain and a loss. The message could be – for example – a political communication or an advert for a product.
6. Design a study to test whether loss aversion is acquired through learning.
7. Evaluate two alternative approaches to utility theory and prospect theory.

NOTES

1. The expected value is $(1/2 \times 1) + (1/4 \times 2) + (1/8 \times 4) + (1/16 \times 8) + \ldots\ldots = \infty$
2. Readers who refer back to the early writings on prospect theory will spot that visual representations of the probability weighting function are not a reverse S-shape, but more of a simple curve. However, an in-depth study by Gonzalez and Wu (1999) indicates that reverse S-shape better captures the weighting function.
3. Even more specifically, the main areas involved were the Right-dorsolateral prefrontal cortex, the Right-intraparietal sulcus, and the Left-intraparietal sulcus.
4. Dopamine is a type of neurotransmitter, that is, a chemical 'messenger' that passes from one nerve cell (neuron) to another.
5. The certainty equivalent is the amount of money that a person would be willing to accept for certain rather than gamble. In this case, a person's certainty equivalent is the amount of money she would be willing to accept rather than obtain $75 only if the relevant National Basketball Association event comes about.
6. Try not to confuse Ernst Weber, the discoverer of Weber's Law, with Elke Weber, whose research is referred to in this section!

RECOMMENDED READING

Montague, R. (2006). *Why choose this book? How we make decisions.* London: Dutton. A very readable account of how expectations and learned values are represented in the brain.