Decision making

Maarten Speekenbrink and David R. Shanks

Division of Psychology and Language Sciences

University College London

Abstract

This chapter reviews normative and descriptive aspects of decision making. Expected Utility Theory (EUT), the dominant normative theory of decision making, is often thought to provide a relatively poor description of how people actually make decisions. Prospect Theory has been proposed as a more descriptively valid alternative. The failure of EUT seems at least partly due to the fact that people's preferences are often unstable and subject to various influences from the method of elicitation, decision context, and goals. In novel situations, people need to infer their preferences from various cues such as the context and their memories and emotions. Through repeated experience with particular decisions and their outcomes, these inferences can become more stable, resulting in behavior that is more consistent with EUT.

Keywords

Decision making, preference, expected utility theory, prospect theory, multiattribute decisions, decisions from experience

Introduction

Should you take the car, bus, or train to work? What should you eat for lunch? Should you join a pension fund or make your own arrangements for retirement? A day is filled with decisions, from mundane to profound. The scientific study of decision making concerns both how people ought to, and how they actually do make such decisions. It is an interdisciplinary field, with contributions from (in alphabetical order) economics, mathematics, philosophy, psychology, and statistics.

In this chapter, we will provide an introduction to decision making research. Given the breadth of the field, this overview will necessarily be sketchy at times. Our aim is to provide the reader with key principles and results and to show where (we think) the field is heading. Keeping with tradition, we start with an overview of (subjective) expected utility theory, the dominant normative theory of decision-making. We then discuss problems with this theory as a descriptive account of decision-making and present a proposed solution, Prospect Theory. In the remainder of the chapter, we discuss evidence that preferences are often unstable and subject to influence from the immediate context and currently active goals. Unstable preferences are mainly found in situations in which decision makers have little experience, where they have to infer their preferences from the available information, such as the immediate context or their own memories and emotions. We finish with an overview of recent research that shows that, when people are allowed to repeatedly make decisions and experience their outcomes, preferences can become stable and people can learn to behave in accordance with the principles of EUT.

How to make decisions: Expected Utility Theory (EUT)

Normative theories of decision making are concerned with how rational people *ought* to make decisions. The most widely accepted normative principle of decision making is the maximization of expected utility. In expected utility theory (EUT), decision problems are analyzed in terms of acts, states, and consequences. Acts are the courses of action a decision maker can follow; they are basic units between which the decision maker can choose and are thus under the control of the decision maker. Each act can have a number of potential consequences and every consequence has some value to the decision maker. Which consequence will actually follow an act depends upon factors outside the control of the decision maker, collectively referred to as “states of nature”. While the true state of nature is usually unknown, the decision maker is assumed to have some idea about the probability that states will occur. These probabilities can either be objective (“decisions under risk”) or subjective (“decisions under uncertainty”). In either case, the main idea is that the decision maker should choose that act which is expected to provide the most value, or “utility”. The expected utility of an act is a weighted sum of the utilities of its potential consequences, where each weight is the probability that that consequence will occur as a result of the act.

As an example, consider a doctor who is visited by a patient complaining of stomach cramps. Based on this symptom, the doctor thinks the patient either has contracted a virus or suffers from a rather serious case of indigestion. Two strains of the virus are going round: strain A is very serious and if left untreated, will result in severe disability; strain B is relatively mild and if untreated, the patient will likely feel ill for a week, but get better thereafter. The two strains need different treatments; both involve administration of a medicine with side effects, but the side effects for the treatment of strain A are much more serious than those for the treatment of B. If the patient has indigestion, no treatment is necessary, and the doctor can send him home without serious consequences.

In this example, there are three acts (treatment for virus A, treatment for virus B, and no treatment), three states (virus A, virus B, and indigestion), and nine possible outcomes. The acts (a), states (s), and outcomes (o) can be entered into a decision matrix as follows:

	s_1 (strain A)	s_2 (strain B)	s_3 (indigestion)
a_1 (medicine A)	o_{11} (major side effects; u=-20)	o_{12} (ill and major side effects; u=-30)	o_{13} (major side effects; u=-20)
a_2 (medicine B)	o_{21} (disability and minor side effects; u=-100)	o_{22} (minor side effects; u=-10)	o_{23} (minor side effects; u=-10)
a_3 (send home)	o_{31} (disability; u=-95)	o_{32} (ill; u=-5)	o_{33} (fine; u=0)

The decision problem can be viewed as a game in which Nature chooses a state and the doctor an act. What is the doctor's best strategy? This will depend on the doctor's belief regarding the likelihood that Nature chooses each state, and the value of each of the outcomes. To determine the probability of the states, the doctor consulted a recently conducted study, which stated the following probabilities for patients with stomach cramps: P("virus A") =.05, P("virus B") = .1, and P("indigestion") = .85. Although neither virus is very likely, it might be unwise for the doctor to send the patient home, as this could result in severe disability if the patient did happen to have contracted virus A. According to EUT, the doctor should, for each course of action, weight the potential outcomes by their likelihood of occurrence and determine an expectation regarding the outcome of the act. In order to do so, each outcome has to be assigned a numerical value referred to as its *utility*. Example values are given in the matrix above. The expected utility of an act is then computed as

$$EU(a_i) = \sum_{j=1}^{N} P(o_{ij}) \times u(o_{ij})$$

where $P(o_{ij})$ is the probability that outcome o_{ij} occurs and $u(o_{ij})$ refers to its utility. In decisions under risk, the conditional probabilities of the outcome are given, while they are subjectively determined in decisions under uncertainty. The best decision is then the act with the highest expected utility (for these example values, the expected utilities are -21, -14.5 and -5.25 for the three acts respectively, so the doctor should decide to send the patient home).

The principle of maximum expected utility was first proposed by Daniel Bernoulli (1952[1738]) as a solution to the St Petersburg paradox. Previously, decision making was defined in terms of maximizing expected (monetary) value, but the St Petersburg paradox showed that people are only willing to pay relatively small amounts to play a game of chance in which the expected winnings are infinite[1]. Bernouilli's solution was based on the idea that the value of a monetary gain is not the same for everyone: 10 dollars means more to a pauper than to a millionaire. In other words, the subjective value (utility) of money has decreasing marginal returns. By proposing that the utility function for money is concave, Bernouilli was able to give a first account for “risk aversion”, the finding that a sure win is usually preferred to a gamble with the same expected value.

The axioms of EUT

The principle of maximizing expected utility can be justified on the grounds that its consistent application will guarantee that the decision maker obtains the maximum utility “in the long run” (after an infinite number of independent repetitions), but this justification is not to everyone's liking (e.g., Lopes, 1981; Samuelson, 1963). Indeed, it was not until Von Neumann and Morgenstern (1947), in the second edition of their famous book “Theory of games and economic behavior”, proved that maximizing expected utility uniquely satisfies a set of reasonable a priori axioms that EUT became the cornerstone of rational decision making. The Von Neumann and Morgenstern framework considers decisions between gambles or lotteries. A lottery consists of a set of mutually exclusive outcomes, each with an objective and known probability, such that the probabilities sum to 1 for all outcomes (i.e., the outcomes are exhaustive). Later, Savage (1954) extended their framework from such decisions under risk to decisions under uncertainty, showing that decisions which satisfy a set of axioms are made as if they maximize *subjective* expected utility, taking the expectation with respect to a subjective probability distribution.

Since the seminal work of Von Neumann and Morgenstern, simplified and alternative axiomatizations of EUT have been proposed (see e.g., Krantz, Luce, Suppes, & Tversky, 1971; Wakker, 1989). The differences between these are not important for our purposes. The four important axioms are:

1. *Completeness.* This axiom concerns the existence of a preference relation for all options. More formally, given two options (lotteries) *A* and *B*, either *A* is preferred to *B*, *B* is preferred to *A*, or the decision-maker is indifferent between *A* and *B*.

2. *Transitivity.* This axiom concerns the relation between pairwise preferences. More formally, if an option *A* is considered at least as good as option *B*, and option *B* is considered at least as good as option *C*, then option *A* must be considered at least as good as option *C*. The axiom of transitivity guards one against a *money-pump.* Suppose you prefer option *A* to *B*, *B* to *C*, and *C* to *A*. Then you should be willing to pay money to trade option *C* for *B*, pay money to subsequently trade option *B* for *A*, and then to subsequently trade option *A* for *C*, ad infinitum. You would end up without any money, and quite likely with the option you started out with.

3. *Independence*. The axiom of independence states that if option *A* is preferred to option *B*, then the lottery (p:*A*,(1-p):*C*) should be preferred to the lottery (p:*B*,(1-p):C), where (p:*A*,(1-p):*C*) should be read as saying that outcome *A* occurs with probability p and outcome *C* with probability (1-p). As the probability of option *C* is common to both gambles, this should have no effect on their preference. A similar axiom was proposed by

Savage (1954) as the *sure-thing principle*: if someone prefers option *A* when state *S* obtains, but also when state *S* does not obtain, then (s)he should prefer option *A* even when (s)he is uncertain about whether state *S* obtains.

4. *Continuity.* The axiom of continuity states that if *A* is preferred to *B*, and *B* is preferred to *C*, then it must be possible to construct a lottery (*p:A,*(1-*p*):*C*) such that the decision maker is indifferent between this lottery and option *B* for sure.

As already mentioned, this is not the only set of axioms which implies the expected utility principle. An important consequence of all of these axiomatizations is that if someone's decisions respect the axioms, the options can be assigned numerical utilities such that the decisions maximize expected utility. In other words, the decisions can be represented *as if* they were made in accordance to maximum expected utility. In this sense, the tenet of rationality lies within the axioms, not in an explicit process of utility maximization. There is no need to assume that the utilities have any external reality (that they are somewhere "inside" the people making the decisions); the utilities are merely a representation of the choices made. This representational view contrasts with previous notions of utility, which equated it with experienced pleasure and pain (e.g., Bentham, 1948 [1789]).

While the axiomatization of rational decision making seems appealing in its simplicity, requiring only consistency in a set of decisions, it is questionable that such internal consistency always implies rationality. Consider the following example by Sen (1993): Suppose a person at a dinner party is offered the last remaining chocolate on a plate. Out of politeness, she chooses to decline. But if the plate contained two chocolates, she might have decided to take one. Based purely on choices, she seems to prefer nothing to a chocolate in the first situation, but a chocolate to nothing in the second situation. Thus, the choices are intransitive over these two situations, while there is nothing irrational in her decisions in light of her social preferences. One may object that the decision alternatives are not the same in the two situations: in the first, she chooses between the last chocolate or nothing, and in the second situation between the second-last chocolate or nothing. While this objection is sensible, it only strengthens the point that the representation of the decision situation must be in accordance to the decision maker's goals: whether decisions are truly transitive or not may not be immediately obvious.

Descriptive failures of EUT

While generally (though not universally) accepted as the normative theory of decision-making, EUT offers a relatively poor description of how people actually make decisions. Over the years, many empirical results have questioned the descriptive validity of EUT and the axioms underlying it. We will only mention a few key results here (for a more extensive overview, see e.g., Schoemaker, 1982; Kahneman & Tversky, 2000).

Allais paradox

In 1953, Allais presented a major challenge to EUT. He showed that when people are presented with the choice between lotteries

A: Receive $1 million for sure

B: A 10% chance of receiving $5 million and an 89% chance of receiving $1 million (and an implicit 1% chance of receiving nothing)

the large majority prefer A to B. However, when presented with a choice between the lotteries

C: An 11% chance of receiving $1 million

D: A 10% chance of receiving $5 million

the majority prefer D to C. As the two gambles are structurally equivalent between the two situations, this pattern violates the independence axiom[2]. This paradox is characteristic of a general finding referred to as the *common ratio* effect, in which the more risky of two lotteries becomes relatively more attractive when the probability of winning in both lotteries is multiplied by a common ratio.

Ellsberg paradox

In 1961, Ellsberg presented a major problem for subjective EUT (e.g., Savage, 1954) in decisions under uncertainty, where the probability of the outcomes can only be subjectively determined. Suppose an urn is filled with 90 balls. Thirty balls are colored red and the remainder is a mix of black and yellow balls, in an unknown proportion. One ball is randomly drawn from the urn. When people are presented with the choice between lotteries

A: Receive $100 if the ball is red

B: Receive $100 if the ball is black

most people choose A. But when people are presented with lotteries

C: Receive $100 if the ball is red or yellow

D: Receive $100 if the ball is black or yellow

most people choose D. The first result suggests that the subjective probability of drawing a black ball is less than 1/3, yet the second implies it is larger than 1/3[3]. Like the Allais (1953) paradox, the Ellsberg paradox violates the independence axiom as the probability of obtaining a yellow ball is identical in lotteries C and D. According to Ellsberg, people have an aversion to ambiguity, and will avoid it when possible. As the probability of winning is known in gambles A and D, these are preferred.

Framing

While not explicit in the axioms of EUT, it is generally assumed that decisions should respect the principle of invariance (Tversky & Kahneman, 1986; Slovic, 1995), according to which preferences should be invariant to the way in which the options are formulated. This, however, does not seem to hold. For instance, consider the "Asian disease" problem (Tversky & Kahneman, 1981), which concerns an upcoming outbreak of a rare disease expected to kill 600 people. Participants are asked to choose between two proposed programs to combat the disease. Those who are given the options

Program A: 200 people will be saved

Program B: a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved

generally choose program A, while participants presented with the options

Program A': 400 people will die

Program B': a 1/3 probability that nobody will die, and a 2/3 probability that 600 people will die

generally choose program B'. It should be clear that, in terms of lives saved, programs A' and B' are identical to programs A and B respectively. This finding shows that whether a description focuses on gains (people saved) or losses (people died) has a substantial impact on the decisions made. This asymmetry between gains and losses forms an important part of Prospect Theory, to which we turn next.

Prospect Theory

Prospect Theory (PT; Kahneman & Tversky, 1979; Tversky & Kahneman, 1992) was an attempt to make minimal changes to EUT in order to make it descriptively valid (Kahneman & Tversky, 2000). According to PT, decisions involve two distinct stages: an "editing phase" and an "evaluation phase". In the editing phase, people apply various simplifications to the decision problem. For instance, they will remove from consideration dominated options that are worse than others in every respect, as well as outcomes with extremely small probabilities. In the subsequent evaluation stage, the edited options are evaluated in a similar fashion to EUT, as a weighted sum of the subjective value of the outcomes.

A key principle in the evaluation stage is that the utility of an outcome is determined with respect to a neutral reference point, usually the status quo. Traditionally, the utility of monetary values was defined with reference to a person's overall wealth. According to PT, outcomes are evaluated as changes from the reference point and depending on where the reference point is placed, outcomes are conceived as either gains or losses. The difference between gains and losses has a much greater impact on the resulting decisions than a person’s overall final wealth state. This is clearly shown by Kahneman and Tversky (1979), who had one group of participants imagine that they were given $1000. When they were then given the choice between

A: A 50% chance of winning $1000

B: A sure win of $500

most participants (84%) preferred option B. A second group were told to imagine they had been given $2000. When they were subsequently given the choice between

A': A 50% chance of losing $1000

B': A sure loss of $500

most people (69%) preferred option A'. As options A and A’ are equivalent in terms of final wealth (as are B and B’), people who base their decisions on final wealth should make the same decision in both situations. This is not what participants did. Apparently, they did not take the initial bonus of $1000 or $2000 into account, because this aspect was common to the two alternatives. Instead, the choice seemed solely based on the discriminating features, something Kahneman and Tversky refer to as the *isolation effect*. In addition, the results indicate that people react differently to losses and gains. People are risk-seeking for losses, preferring a risky option to a certain option with the same expected value, but risk-averse for gains. To account for this asymmetry, PT proposes a utility function – now referred to as the value function – for gains and losses as depicted in Figure 1. The value function is discontinuous (has a "break") at the reference point and is concave for gains, and convex for losses.

[Figure 1 here]

As shown in Figure 1, the value function also has a steeper slope for losses than wins, which accounts for a second finding, namely that "losses loom larger than gains". This is illustrated in the *endowment effect* (Thaler, 1980). In a classic study showing this effect (Kahneman, Knetsch, & Thaler, 1990), some university students were given a mug and, in a roundabout way, asked for the minimum price they would accept to sell the mug. Another group was not given a mug and asked for the price they would pay to buy the mug. Kahneman et al. found that the median selling price was over twice as large as the median buying price. This difference between Willingness-To-Accept (WTA) and Willingness-To-Pay (WTP) indicates that the (dis)utility of loosing a mug is greater than the utility gained from obtaining it. Another effect illustrating loss aversion is the "status quo bias" (Samuelson & Zeckhauser, 1988), which refers to the tendency to remain in the same state when change involves exchanging a relatively small loss for a larger gain. For instance, Knetsch (1989) randomly gave students either a mug or a chocolate bar. While given the opportunity to trade one for the other, approximately 90% of the students did not do so, showing a preference for whatever they were randomly allocated.

A second key principle in Prospect Theory is the transformation of probabilities by a weighting function. The weighting function reflects the subjective impact of probabilities on decisions. Numerous findings suggest that this impact is not linear (e.g., Camerer & Ho, 1994; Wu & Gonzales, 1996), as it should be according to EUT.

Zeckhauser's Russian roulette example gives an intuitive illustration that the effect of changes in probability is not the same over the whole scale:

> Suppose you are compelled to play Russian roulette, but are given the opportunity to purchase the removal of one bullet from the loaded gun. Would you pay as much to reduce the number of bullets from four to three as you would to reduce the number of bullets from one to zero? (in Kahneman & Tversky, 1979, p. 283)

Most people would pay substantially more to reduce the number of bullets from 1 to 0 than to reduce it from 4 to 3, while the increase in expected utility should be identical in both cases. This result has been called the *certainty effect* (Kahneman & Tversky, 1979): people seem to overweight certain outcomes (in the example, when the number of bullets is zero, it is certain that you will survive) relative to outcomes which are merely probable. PT accounts for this effect, as well as others, by assuming that the probabilities are distorted by a weighting function. Originally, this weighting function was assumed to be identical for gains and losses and applied directly to the probabilities. This, however, led to some counter-intuitive predictions. In a later refinement, called Cumulative Prospect Theory (CPT, Tversky & Kahneman, 1992), the weighting function is applied to cumulative probabilities, as in rank-dependent utility theory (Quiggin, 1982). In CPT, the cumulative probabilities are computed separately for each prospect (decision alternative), and separately for gains and losses within each prospect. For gains, the outcomes are ordered in increasing value, and each outcome is associated with a cumulative probability of receiving that outcome or anything better. A weighting function (see Figure 2) is applied to these cumulative probabilities, and for each option, the decision weight is defined as the difference between the transformed cumulative probability for that gain and next highest outcome. The procedure for losses is similar, but the cumulative probabilities are probabilities of obtaining a loss or anything worse, and the decision weights are computed as the difference between the transformed cumulative probability of an outcome and the immediately preceding loss.

[Figure 2 here]

This cumulative weighting procedure is intricate and perhaps not immediately plausible from a psychological viewpoint. Regardless, Prospect Theory has proven a successful descriptive account of decision making (see e.g., Edwards, 1996, for a review). On the other hand, it does not provide deep psychological insight into many of the processes it proposed (Newell, Lagnado, & Shanks, 2007) and a complete account of decision-making will likely take more factors into account. We will discuss a number of such factors later on, but a prominent one we'll briefly mention now is regret (Loomes & Sugden, 1982). When making a decision under uncertainty, the outcome informs us about the actual state of the world. Suppose that, after deciding not to treat the patient, the doctor in the earlier example finds out that the patient had actually contracted virus B. Had she decided to treat the patient with the appropriate medicine, she would have saved him a week of illness. Knowing this, the doctor now regrets her decision, and the displeasure of this regret is something over and above the disutility of the illness. If, on the other hand, the doctor had found out the patient did indeed suffer from indigestion, she would rejoice at her decision, and this rejoicing adds something over and above the utility of a healthy patient. Loomes and Sugden's (1982) Regret Theory offers an account of decision-making based on the anticipation of such regret and rejoicing. Effectively, regret theory extends EUT by assuming that regret and rejoicing modify the classic "choiceless" utility function, thus making utility context-dependent.

Multi-attribute decisions

Most real-life decisions are not posed as choices between monetary gambles. For instance, when deciding between two job offers, there are many relevant factors besides monetary outcome (salary), such as the convenience of the location, the sociability of your colleagues, and so on. The decision alternatives have values on many different attributes, and each attribute might be relevant to a different goal (e.g., maximizing earnings, minimizing daily commute, extending social network, etc.). Usually, there is a necessary trade-off between such attributes. For instance, a more highly paid job might require a longer commute. The question then arises how these attributes are integrated when making a decision. Traditionally, multi-attribute utility theory has taken the view that people first evaluate each option on every attribute, after which these "single-attribute utilities" are weighted according to the relative importance of the attributes and then summed to give the option's overall utility (e.g., Keeny & Riaffa, 1976). In this view, utility serves as a "common currency" such that a relative lack in one attribute (e.g., salary) can be compensated for by a relative abundance on another attribute (e.g., location). But this compensatory nature of multi-attribute decision making has been questioned. For options with many attributes, the computation of overall utility seems too complex, and researchers have proposed that people use a variety of heuristics to deal with this complexity.

One of the first people to propose that a theory of rational decision making should take people's cognitive limitations into account was Simon (1955, 1956), who coined the term "bounded rationality". According to Simon, people *satisfice* rather than optimize; they look for options that are good enough, rather than the best. In multi-attribute decisions, satisficing consists of setting a minimal value for each of the attributes. As soon as an option is found which surpasses all these lower bounds, a decision is made. Satisficing seems like an economical strategy: options can be searched attribute-wise, and any failing options can be eliminated from the choice set, thus sequentially reducing the complexity of the procedure. There is even research that suggests that while optimizers achieve better outcomes, satisficers are more content with the outcomes they achieve[4] (Iyengar, Wells, & Schwarz, 2006). By considering more options, maximisers may be more prone to anticipate and experience regret (Schwartz et al., 2002).

Tversky's (1972) Elimination-by-Aspects model effectively combines satisficing with optimization. While options that do not meet set criteria are eliminated from the consideration set, this elimination continues until one alternative remains. Elimination-by-Aspects proposes that the search through attribute values is stochastic, with the probability of choosing an attribute dimension depending on its overall importance to the decision-maker. As Tversky shows, this model can account for certain violations of the independence axiom. Elimination-by-Aspects is an example of lexicographic decision-making, in which attributes are considered sequentially and there is no inter-attribute compensation. A more recently proposed lexicographic decision strategy is Take-the-Best (TTB; Gigerenzer & Goldstein, 1996). This decision-making heuristic was originally proposed for inferential decisions in which there is an objectively correct alternative. TTB is similar to Elimination-by-Aspects, but does away with random sampling of attribute dimensions in favor of a fixed search-order according to attribute validity and eliminates all alternatives which do not have the maximum value on a considered attribute. This last aspect makes TTB a very frugal strategy, as a decision is made as soon as an attribute is found on which a single alternative is best. As lexicographic strategies, Elimination-by-Aspects and Take-the-Best are non-compensatory. They can conflict with compensatory strategies such as MAUT, because an alternative that is poor on one attribute, and therefore eliminated by a lexicographic strategy, might be good on all other attributes.

Claims that compensatory strategies are too complex in all but the simplest decision problems have resulted in significant effort to delineate the factors determining decision strategy use. An interesting suggestion is that people adapt their strategies to the decision environment in an attempt to minimize cognitive effort whilst achieving a satisfactory level of decision accuracy (Gigerenzer, Todd, & the ABC Research Group, 1999; Payne, Bettman, & Johnson, 1993; Rieskamp & Otto, 2006). Simulation studies have shown that in inferential tasks strategies like TTB result in surprisingly accurate decisions, sometimes outperforming compensatory strategies such as linear regression (e.g., Gigerenzer & Goldstein, 1996), although the generality of this conclusion has been questioned (Chater, Oaksford, Nakisa, & Redington, 2003). In any case, potential success provides no direct evidence that people actually employ TTB when making decisions. Studies that have investigated this empirical claim directly have found some evidence for TTB usage (e.g., Bergert & Nosofsky, 2007; Bröder, 2000; Newell & Shanks, 2003; Newell, Weston, & Shanks, 2003; Rieskamp & Hoffrage, 1999). While people usually search information in order of cue validity, they often don't stop their search after the first discriminating cue has been found (e.g., Newell, Weston & Shanks, 2003). Although simple rules like TTB are employed some of the time, studies show large variability, both within and between people. The growing consensus seems to be that people can use both compensatory and non-compensatory strategies, and adapt their strategies to the requirements of the task (Bröder & Newell, 2008; Gigerenzer & Gaissmaier, 2011; Rieskamp & Otto, 2006).

The instability of preferences

The traditional economic view is that decisions reveal someone's stable preferences. You either prefer the Rolling Stones to the Beatles, or you don’t. According to a rather naïve view, people make decisions by reading the utility of the various outcomes from a lookup table. Access to such a stable preference relation implies the principle of invariance (Tversky & Kahneman, 1986; Slovic, 1995), according to which preferences should not depend on the description of the outcomes (description invariance), nor on the method of elicitation (procedure invariance). But there is considerable evidence that neither of these holds universally.

Procedure variance

Evidence against description invariance has already been mentioned (e.g., the Asian disease problem). A classic study that showed procedure variance was conducted by Lichtenstein and Slovic (1971). They presented participants with the following lotteries:

P: A 95% chance of winning $2.50 and a 5% chance of losing $0.75

D: A 40% chance of winning $8.50 and a 60% chance of losing $1.50

In the P ("probability") lottery, there is a relatively large probability of winning a modest amount, while in the D ("dollar") lottery, there is a modest probability to win a relatively large amount. When asked which gamble they would prefer to play, the majority of the participants chose the P lottery. However, when participants were subsequently asked to give their minimum selling price for a ticket in the lottery, lottery D was usually (in 88% of the cases) given a higher price than lottery P (implying a preference for D).

Preference reversals have subsequently been found when comparing other preference elicitation procedures, such as choice and matching (Slovic, 1975; Tversky, Sattath, & Slovic, 1988), and separate and simultaneous evaluation of options (Hsee, 1996, List, 2002). The study by List (2002) is particularly relevant, as it shows that preference reversals also occur outside the laboratory. List studied people buying baseball cards at a specialist convention. In one condition, people placed bids on two sets of baseball cards viewed simultaneously. One set contained 10 cards, all in excellent condition, while the second set contained three additional cards in very poor condition. As expected, people generally placed somewhat higher bids on the 13 card bundle. However, when people in a second condition evaluated the sets separately (either bidding for the 10 card or the 13 card bundle), the bids for the set of 10 cards were higher than those for the 13 card set. According to Hsee (1996), such effects depend on the "evaluability" of the attributes. Some attributes, like the number of baseball cards in a set, may be difficult to evaluate in isolation; their worth depends on what's offered in other options. Because of this indeterminacy, these attributes will have little influence on separate evaluations. But due to its inherently comparative nature, they will inform joint evaluation.

Context effects

The above finding suggests that the decision alternatives can provide a meaningful context to ground preferences. There are other key findings supporting the general context dependence of preferences. Three main effects are the similarity, compromise, and attraction effects. These effects all consider how the choice probabilities between two options are affected by the addition of a third alternative and all involve a violation of the axiom of independence.

[Figure 3 here]

A schematic representation of the contexts in which these effects are found is given in Figure 3. The two-option comparison always concerns options A and B. To describe the effects, suppose that in a binary choice between these two options, (a group of) people show indifference between the options, so that the proportion choosing A equals the proportion choosing B, i.e., $P(\text{A}|\{\text{A,B}\}) = P(\text{B}|\{\text{A,B}\})$.

The *similarity effect* (Tversky, 1972) refers to the finding that the addition of option S (which is similar to option A, but slightly worse on attribute 1 and slightly better on attribute 2) to choice set {A,B} mainly takes away from the probability of choosing A, i.e. $P(A|\{A,B,S\}) < P(A|\{A,B\})$, whilst leaving the probability of choosing B relatively intact. As a result, people can prefer A over B in the {A,B} context, but B over A in the {A,B,S} context, which violates the independence axiom. As argued by Tversky (1972), due to the similarity between options A and S, these may be taken as roughly equivalent, and a choice may be conducted in two stages: in the first stage, people choose between option B or the set A and S. If the set is chosen, then the second stage involves the choice between A and S. Indeed, Tversky's (1972) Elimination-by-Aspects model was explicitly formulated to account for the similarity effect.

The *attraction effect* (Huber, Payne, & Puto, 1982; Simonson, 1989**)** refers to the finding that the addition of option D (which is similar to A, but worse on both attributes) to the choice set {A,B} increases the number of people choosing A, i.e. $P(A|\{A,B,D\}) > P(A|\{A,B\})$. A marketing example of the effect is when a manufacturer introduces a new product (D) onto the market which is clearly inferior to one they already produce (A), in the hope of taking market share away from a competitor brand (B). The key difference between the situation leading to the similarity effect and the attraction effect is that option D is asymmetrically dominated (it is worse than A on both attributes, but worse than B on only one attribute), while option S is not dominated (i.e., it is better than A on one attribute). The attraction effect violates the *regularity principle* (Tversky, 1972), according to which the addition of an option to the choice set can never increase the probability of choosing an option relative to the original set (i.e., the principle implies that $P(A|\{A,B\}) \geq P(A|\{A,B,D\})$). The regularity principle is a weaker form of the independence axiom and is implied by many decision-making models, including

Elimination-by-Aspects and a large class of random utility models. Therefore, the attraction effect poses a major problem for many accounts of decision-making.

Another violation of the regularity principle is found in the *compromise effect* (Simonson, 1989; Tversky & Simonson, 1993). Suppose that people are indifferent between the options in a series of binary choices between pairs {A,B}, {A,C} and {B,C}, i.e. P(A|{A,B}) = P(B|{B,C}) = P(C|{A,C} = .5). The compromise effect refers to the finding that in a choice from the set {A,B,C}, people often show a clear preference for the compromise option C, i.e., $P(\text{C}|\{\text{A,B,C}\}) > P(\text{C}|\{\text{A,C}\}) = P(\text{C}|\{\text{B,C}\})$.

Simonson (1989) proposed that the attraction and compromise effects arise because people find it easier to justify their choices in the enlarged choice sets. For instance, the dominated option D provides a clear justification for choosing A. The importance of such justifications is highlighted in the reason-based choice account of Shafir, Simonson, & Tversky (1993), which shows how various choice anomalies can be more readily understood when considering how the decision context can affect the generation of reasons for and against choosing various alternatives. A more tractable value, rather than reason, based account of the attraction and compromise effects was given by Tversky and Simonson (1993) in their context-dependent advantage model. However, this model cannot account for the similarity effect (Roe, Busemeyer, & Townsend, 2001). A theory that can account for all three context effects is Decision Field Theory (Busemeyer & Townsend, 1993; Roe et al., 2001), to which we'll return later on.

Goals

Outcomes with multiple attributes are usually relevant to a number of goals. As people generally have multiple goals at a single time, multi-attribute decisions will usually involve a trade-off between the attainment of different goals (Payne, Bettman, & Johnson, 1992). When multiple goals conflict, selective attention to different subsets of these goals can influence both how a decision is made and the outcome of this process (Krantz & Kunreuther, 2007; Markman & Medin, 2002). If there is a shift in the relevance of different goals between one decision and the next, decision-makers can seemingly make inconsistent choices even though these decisions are individually all instrumentally rational.

Brendl, Markman, and Messner (2003) conducted a study which shows how goal activation can affect decisions. They offered habitual smokers the opportunity to buy tickets for a lottery involving either a cash or cigarette prize. Each lottery would be conducted after two weeks, so neither prize could be used to satisfy current needs. The crucial manipulation was in the timing of the offer: half the smokers were asked before and half after having the opportunity for a post-class cigarette. The results showed that those who had not smoked yet bought significantly fewer tickets to win money than those in the other group. The authors also found such devaluation for options irrelevant to currently active goals when hungry participants rated the attractiveness of non-food items (see also Markman, Brendl, & Kim, 2007).

Maintaining a set of conflicting goals may be subserved by a process called mental accounting (Markman & Medin, 2002; Thaler, 1985). Mental accounting is often considered a hindrance to rational decision making, but it may actually help people to obtain long-term goals in the face of currently active short-term goals (Shefrin & Thaler, 1992). An example of mental accounting comes from Kahneman and Tversky (1984), who found that people were less likely to buy a theatre ticket they previously lost compared to people who had not previously bought a ticket but lost an amount of money equal to the ticket price. In contrast to the view that money is “fungible” (substitutable), this finding suggests that people segregate theater costs from other costs. Those who pre-purchased a ticket were less willing to spend additional money from the theater account, while the money lost by people in the other condition could have come from any other account. Although this behavior may seem unreasonable (money lost is money lost), associating different goals with separate accounts (e.g., entertainment, education, etc.) may help people protect particular goals from other, more immediately pressing ones.

The inference of preferences

In cases where decision-makers have little prior experience with the options, their preferences may not be immediately obvious and they will need to infer the subjective value of the outcomes. Indeed this may even be true of some decisions made repeatedly. In the literature, this process is often called preference construction (e.g., Payne et al., 1992; Slovic, 1995; Tversky & Thaler, 1990), but we prefer the term preference inference, as we believe this reflects the aims and constraints of the decision makers better. While preference construction implies the absence of inherent preferences, preference inference implies that underlying preferences may exist, but that these need to be discovered (cf. Plott, 1996; Simonson, 2008). In attempting to infer their preferences, people may rely on a variety of cues, such as those in the immediate decision context, memory of previous experiences, and affective reactions.

Context effects (again)

One notable finding is that the range of attribute values can have a large effect on how they are perceived. For instance, Mellers and Cooke (1994) had students rate the attractiveness of various apartments which differed in rent and distance to campus. Mellers and Cooke found that the same change on an attribute (e.g., a rent change from $200 to $400) had a greater effect on perceived attractiveness when the range of rents considered was narrow (e.g., all apartments had rents varying between $200 and $400) than when the range was wide (e.g., rent varied from $100 to $1000). Indeed, this effect was so pronounced that it resulted in preference reversals. For instance, a $200 apartment 26 minutes from campus was preferred to a $400 apartment 10 minutes from campus when the range of rents was small but the range of distances was wide, while the opposite was found when the range of rents was wide and the range of distances was narrow. Mellers and Cooke also showed that this effect was not due to a change in relative attribute weight; rather, the perceived value on the attribute differed as a function of the range of values on that attribute. Although the students were presumably quite familiar with the decision context, they seemed to rely on relative rather than absolute attribute values in their preferences. This makes sense if you consider that assessments such as whether an apartment is expensive are inherently relative, as are many other qualities, such as sport achievements and what constitutes a well-paid job. Especially in relatively unknown environments, the distribution of attribute values can provide useful information, if only about what others conceive as important. In support of this, Beattie and Baron (1991) have found that range effects on relative attribute weights depend on participants' experience with a scale.

Memory

It seems plausible to assume that, in order to estimate the value of various decision outcomes, people will attempt to retrieve from memory prior experiences with similar outcomes. While it is hard to deny such a role of memory in evaluation, there is relatively little work that directly assesses the effect of memory on decision making (Weber & Johnson, 2006).

Stewart, Chater, and Brown (2006) have formulated a theory in which memory plays a central role. According to their Decision by Sampling (DbS) theory, people value decision alternatives by evaluating their attributes against a sample of other attribute values. This decision sample consists of attribute values in the immediate decision context and values from memory of previously encountered attributes. People are assumed to only make ordinal comparisons, and the value of a decision alternative is determined by its relative ranking in the sample. They show that Prospect Theory's value function closely matches relative rankings of credit and debit amounts in actual bank transfers. Assuming that memory provides a veridical representation of the environment, DbS thus provides an explanation for the shape of the value function. In a similar way, it provides an account of the probability weighting function, and hyperbolic time-discounting.

DbS assumes, at least as a first approximation, that attributes are sampled randomly from memory. This assumption is unlikely to hold, given what is known about memory. For instance, memory is associative, and retrieved items are likely to be similar to previously retrieved items. In addition, there are interference effects in memory, such that retrieved items can inhibit the retrieval of other items (e.g., Anderson & Spellman, 1995). Johnson, Häubl, and Keinan (2007) argue that such memory interference may play a key role in the endowment effect. According to their Query Theory, people first consider reasons for maintaining the status quo, and then consider reasons for changing it. Reasons are considered to be constructed by queries of memory, and due to memory interference, query order can have strong effects on the outcome of this process. Because sellers first consider reasons for keeping the mug, which enhance its value, retrieval of value-decreasing reasons for selling the mug is inhibited, resulting in a more positive evaluation of the mug compared to buyers who first consider reasons for keeping their money (value decreasing) before reasons for buying the mug (value enhancing). In an experiment, Johnson et al. found some support for this account. In particular, when the assumed natural order of queries was reversed (sellers were guided to give value decreasing aspects before value increasing aspects, and vice versa for buyers), the endowment effect was eliminated.

Being able to retrieve any memories for an alternative may in itself be value enhancing. Preferences for products and other objects tend to be related to their familiarity (e.g., Hoyer & Brown, 1990). In this way, recognition can serve as a cue to value. According to Goldstein and Gigerenzer (2002), recognition is a powerful heuristic in decision making. As recognition is often related to quantities of interest, such as the quality of higher education institutes, or a team's success in sports, basing decisions on recognition can give good outcomes. A telling example is in stockmarket investment, where it has been shown that portfolios comprised of the most-recognized options, on average, outperformed investment experts and managed funds such as the Fidelity Growth Fund (Ortmann, Gigerenzer, Borges, & Goldstein, 2008). According to the recognition heuristic, if only one decision alternative is recognized, that alternative is chosen and no further information is used. While most agree that recognition is a useful and often informative cue, this last claim has led to some debate (e.g., Oeusoonthornwattana & Shanks, 2010; Pachur, Bröder, & Marewski, 2008). For instance, Oeusoonthornwattana and Shanks (2010) showed that when, in addition to recognition, other information about options is available, people will not ignore this. As such, recognition may be a cue to value, but it is not used exclusively.

While people may explicitly attempt to retrieve experiences in order to assess the value of decision alternatives, memory effects may be more subtle. As shown in a study by North, Hargreaves, and McKendrick (1997), people can be primed to make certain decisions. North et al. investigated wine-buying behavior in a supermarket where four French and four German wines were displayed. When French background music was played, French wines outsold the German wines, while the reverse was true when

German music was played. Questionnaires indicated that the music made people think of its originating country, thereby retrieving more memories relating to a particular wine region. However, most people denied the music influenced their wine choice. Additional support for priming effects in decision-making comes from Berger, Meredith, and Wheeler (2008), who report evidence that the location where people vote can influence their decision. When people were assigned a school as polling station, they were more likely to support a school funding initiative than when they voted in other locations. According to the authors, the context may have primed people to retrieve favorable arguments for the school funding initiative.

By retrieving previous experiences with outcomes, people should be able to make an informed evaluation of the decision alternatives. However, evidence suggests that "remembered utility" is not a veridical representation of directly experienced (dis-)pleasure (Kahneman, Wakker, & Sarin, 1997). In a classic study, Redelmeier and Kahneman (1996) asked patients undergoing a painful colonoscopy to rate their level of discomfort every 60 seconds. After the procedure, they were asked to rate their overall discomfort. These final ratings did not reflect an integration of the minute-by-minute experienced discomfort. Rather, total ratings seemed to be formed by a "peak-end" rule, averaging the maximum discomfort and the discomfort experienced just before finishing the procedure. For instance, a patient whose procedure ended with a period of relatively mild discomfort rated the overall discomfort as less than another patient whose procedure ended abruptly at a point of relatively large discomfort, even though the first patient's procedure lasted considerably longer and he thus accumulated more overall discomfort. Similar results were obtained in a study in which participants watched pleasant and unpleasant movies (Fredrickson & Kahneman, 1993).

As shown in the studies by Kahneman and colleagues, we do not always remember events in a way that accurately reflects how they were experienced. Our recollections can distort past events and how enjoyable or unpleasant they were. Therefore, experienced utility is not always a good predictor of future decisions. This is especially clear in a study by Wirtz, Kruger, Scollon, and Diener (2003), who showed that students' desire to repeat a type of vacation was better predicted by their recollected than their actually experienced enjoyment. Although recollections were related to experienced enjoyment, there was also an independent effect of predicted enjoyment before the vacation. Expectations thus seem to have a long lasting effect which is not overwritten by actual experience. As recollected enjoyment mediates the effect of actual experience, when determining whether someone will repeat an experience such as revisiting a restaurant, asking them during the experience will be less useful than seeking their subsequent remembered experience.

Emotion

While the traditional revealed preference view on decision making eschewed direct consideration of the hedonic value of decision outcomes, recent work has begun to explore the role of emotion in decision-making. Perhaps the most extensive thesis on the role of emotion in decision making is Damasio's (1994, 1996) Somatic Marker Hypothesis. According to Damasio, in complex and conflicting decision problems, people rely on emotion-based biasing signals generated from the body that simplify the problem by marking risky and potentially costly alternatives. Evidence for this Somatic Marker Hypothesis comes from patients with damage to the ventromedial prefrontal cortex (VMPFC), such as the famous Phineas Gage (see Damasio, Grabowski, Frank, Galaburda, & Damasio, 1994), who survived an accident in which an iron rod went through his head and destroyed much of his frontal cortex. While Gage seemed to have no intellectual impairment, he became impulsive and unpredictable and started showing strange social and decision-making behavior, consistent with the loss of his somatic marker system.

Experimental evidence for the Somatic Marker Hypothesis has mainly relied on the Iowa Gambling Task (IGT; Bechara, Damasio, Damasio, & Anderson, 1994), which was designed to simulate real-life decision making in the way it factors uncertainty, reward, and punishment. The IGT is a multi-armed bandit task in which participants repeatedly choose to draw a card from four decks. Each card will provide a (monetary) reward, but some cards also provide an additional punishment. Two decks are disadvantageous: whilst providing higher rewards, they are also associated with higher punishments, and consistently choosing these decks will result in an overall net loss. The other two decks are advantageous: while providing lower rewards than the disadvantageous decks, the punishments are also lower and consistently choosing these decks will result in an overall net gain. Bechara et al. (1994) found that people with damage to the VMPFC were relatively impaired in this task (choosing more from the disadvantageous decks than healthy controls and patients with other brain lesions) and concluded that the VMPFC plays a key role in the emotional evaluation of decision outcomes. In further support of the Somatic Marker Hypothesis, later studies found that healthy participants, but not those with damage to the VMPFC, developed anticipatory skin conductance responses (SCRs) to choices from the disadvantageous decks (Bechara, Tranel, Damasio, & Damasio, 1996; Bechara, Damasio, Tranel, & Damasio, 1997). These anticipatory reactions correlated with improved performance and apparently developed before participants acquired explicit knowledge about the structure of the task (Bechara et al., 1997). While this suggests that non-conscious emotional signals, or somatic markers, guide normal decision making before conscious knowledge does, this conclusion was later placed in doubt by Maia and McClelland (2004). Using more sensitive questions, they showed that participants developed awareness of the good and bad decks much earlier than previously thought. While the many studies using the IGT have failed to provide compelling support for the Somatic Marker Hypothesis (for a review, see Dunn, Dalgleish, & Lawrence, 2006), there seems little doubt that emotions play a significant role in decision making.

The role of emotional reactions to outcomes as a guide to decision-making has been formulated as the *affect heuristic* (e.g., Finucane, Peters, & Slovic, 2003). Decision problems can conjure vivid images with attached affective reactions, which influence the overall evaluation of alternatives. Some evidence for this idea was gathered by Finucane, Alhakami, Slovic, and Johnson (2000), who showed that they could manipulate the perceived benefit of nuclear power by telling participants that the associated risk was either high or low. The authors argued that high risk induces an overall negative affective evaluation, which, via an affect heuristic, results in an inference of low benefit. In a similar way, perceived risk could be manipulated by informing participants that the benefits were either high or low. According to the “risk as feelings” hypothesis (Loewenstein, Weber, Hsee, & Welch, 2001), such affective reactions can also explain the overweighting of small probabilities discussed previously. According to this account, when moving from impossibility (i.e., $p = 0$) to possibility (e.g., $p = .01$), a threshold is crossed in which a consequence of no concern becomes a source of worry or hope. As possible consequences are imagined, associated emotional reactions arise, but this process is assumed to be little affected by further increases in probability. In the words of Loewenstein et al. (2001, p. 276) “One's mental image of what it would be like to win the state lottery [...] is likely to be about the same, whether there is a 1 in 10,000,000 chance of winning or a 1 in 10,000 chance of winning.” The idea that departures from certainty are especially marked for affect-rich compared to affect-poor consequences was supported by Rottenstreich and Hsee (2001), who found a preference reversal between certain and possible outcomes. In their study, participants chose between a lottery with a 1% chance to win a kiss from their favorite movie star (affect-rich) or a lottery with a 1% chance to win $5. The large majority preferred the kiss lottery, even though when participants were given the choice between the outcomes for certain, the large majority preferred the $5 prize. Rottenstreich and Hsee explain this reversal by assuming that the nonlinearity of the probability weighting (the "S-shapedness" in Fig. 2) increases for affect-rich outcomes.

The role of emotion in decision making is receiving more and more attention (Weber & Johnson, 2009). It is increasingly common to find dual-process accounts of decision-making, which contrast intuitive and analytical decision-making processes (e.g., Kahneman & Frederick, 2002; Sloman, 1996). Intuitive, "System 1" processes are automatic, effortless, and associative, and work on affective reactions and specific memory exemplars. Analytical, "System 2" processes are deliberate, effortful and deductive and work on abstract, affectively neutral inputs. System 1 processes are assumed to propose intuitive solutions to decision problems. These proposals are monitored by System 2 processes, which may endorse, correct, or override the inputs from System 1 processes (Kahneman & Frederick, 2002). Although the claim of qualitatively separable processing systems is controversial (e.g., Keren & Schul, 2008), it seems plausible that various cognitive and affective processes provide cues to solve a decision problem. This cue generation may appear more or less automatic. Integrating these cues into a final decision may be more akin to what is understood by "System 2" processes. As discussed earlier, there are different strategies for multi-attribute decisions and whether cues are integrated in a compensatory manner may depend on the time available and the effort one is willing to put in.

Predicting future preferences

As noted by March (1978), rational decision-making involves two types of inference: "guesses about future consequences of current actions and guesses about future preferences for those consequences" (p. 40). While (subjective) EUT assumes a completely specified set of potential consequences and their subjective likelihood of occurrence, as well as a complete preference order over them, it seems more realistic to assume that these are only vaguely specified. As current preferences for outcomes may only be weak indicators of future preferences for those outcomes, preference prediction is an inherently uncertain endeavor. For example, when deciding to buy a house, you need not only consider what it would be like to live there immediately, but also how you will feel in a few years' time. The difficulty with such prediction is that your needs and preferences are likely to change over time. For instance, an apartment in the heart of the entertainment district may be very satisfactory for a young couple, but may not be suitable after they have their first child. But even in the absence of such major life changes, people seem to have problems in predicting their future preferences.

Studies on affective forecasting have shown that people are generally poor at predicting their affective reactions to events in the future (Gilbert, Gill, & Wilson, 1998). Perhaps this is not too surprising, as accurate foresight is a desirable but illusive cognitive faculty. While evidence suggests that people are reasonably accurate when predicting their immediate reactions to outcomes, they have little insight into how their preferences will adapt over time (Loewenstein, O'Donoghue, and Rabin, 2003). A common finding is that predictions of future utility are anchored on current utility. For instance, hungry shoppers tend to buy more than intended (e.g., Gilbert, Gill, & Wilson, 1998), apparently expecting to be as hungry in the future. This unwarranted projection of a current state into the future has been called the *projection bias* (Loewenstein et al, 2003). While current liking of options can be indicative of future liking - people are quite adept at predicting whether future events will be pleasant or unpleasant - people seem relatively poor at predicting the future intensity of their feelings (Wilson & Gilbert, 2005). For instance, people often overestimate the impact of life-changing events, such as acquiring a medical condition, or a relationship breakup. While people readily adapt to new circumstances, their predictions seem mostly based on the initial impact of the event (Wilson & Gilbert, 2005).

An early study by Kahneman and Snell (1992) illustrates how people fail to predict their changing preferences in a more mundane situation: tasting yoghurt. Their participants were unable to predict that repeated exposure to plain yoghurt, rather than decreasing, would actually increase their liking. Overall, there was little correlation between participants' predictions and actual liking. In a study by Simonson (1990), participants who chose three snacks in advance for consumption on later occasions made more varied choices than those who chose their snacks on those occasions for immediate consumption. This indicates that participants may have wrongly assumed their tastes would vary. Interestingly, when participants were asked, after an initial choice, to predict their choices on consecutive occasions, the difference between simultaneous and separate choice disappeared. While Simonson found that people’s preferences are more stable than they expected, the results of Kahneman and Snell’s study suggest that people should expect preference change, but that they have difficulty predicting its direction. That people seem to have little insight into how their preferences will change, even for the very near future, was also shown by Loewenstein and Adler (1995). In their experiment, people were unable to predict how endowment would change their preferences.

The process of preference inference

If preferences are not simply accessed but inferred, this process should be of interest to decision-making researchers. A common view is that this process involves a form of sequential sampling. Sequential sampling models of decision making have a long history in psychology (e.g., Audley, 1960; Ratcliff, 1978). A common theme is that, during the course of deliberation, evidence in favor of decision alternatives is accumulated until the decision maker is satisfied (s)he has a good-enough basis for making a final choice. Evidence is assumed to be randomly sampled on a moment-by-moment basis (hence the name sequential sampling), resulting in a stochastic process of evidence accumulation. In "race" models, evidence for each alternative is accumulated separately. In "diffusion" models, accumulation involves the relative evidence in favor of one option over others. In both cases, a decision is made when the accumulated evidence (whether absolute or relative) exceeds a threshold. The setting of the threshold is important for speed-accuracy trade-off: setting a lenient criterion will result in quick but unreliable decisions, while setting a stricter criterion will result in a more reliable but lengthy decision-making process.

An extensive overview of the various sequential sampling models exceeds the scope of this chapter (the interested reader may consult e.g., Smith, 2000). We will only describe one of them in more detail here, namely Decision Field Theory (DFT; Busemeyer & Townsend, 1993; Roe, Busemeyer, & Townsend, 2001). DFT assumes a person deliberates over the decision alternatives by thinking about the consequences of each action. During deliberation, attention shifts between the possible consequences, giving rise to affective reactions which are accumulated in an overall preference state for each action. Affective reactions to consequences can be positive or negative. In DFT, these affective reactions result in a valence for each option, which reflects the relative advantage of an option over the other options. Valences are accumulated into an integrated preference state for each option. Crucially, this integration process is competitive and incorporates an inhibitory mechanism which increases the advantage of a preferred option over less preferred, similar options. An example of the dynamic trajectories resulting from this integrative process is given in Figure 4. As in other sequential sampling models, a decision is made as soon as a trajectory surpasses a decision threshold.

[Figure 4]

As already mentioned, Decision Field Theory can account for the similarity, attraction, and compromise effects. The similarity effect results from the attention switching process. When considering options S, A, and B (Figure 2), if attention is focused on dimension 1, option B receives a large advantage over S and A. But if attention is focused on dimension 2, both S and A have a positive valence, while option B has a negative valence. Therefore, addition of option S to the choice set takes away from option A, but not option B. The attraction affect results primarily from the competitive integration process. Because option D is similar to option A, they compete more than options D and B, and this results in a relative advantage of option A over D. The compromise effect depends on both attention and inhibition, but the precise explanation for this effect is more intricate (see Roe et al., 2001, for more details). Another sequential sampling model that can account for all three effects is the Leaky Competing Accumulator model (LCA; Usher & McClelland, 2001; Usher & McClelland, 2004). While the relative merits of DFT and the LCA, and other models like them, needs further research, sequential sampling models can offer a precise theoretical framework to synthesize a wide variety of empirical results on decision making.

Learning to make decisions

Much of decision research has concerned how people make one-shot decisions between fully specified gambles, or otherwise abstracted choice situations. But how relevant are the results for everyday decision-making outside the laboratory? How often does one face a never-to-be-repeated situation in which the options are completely and accurately described? Given relatively little experience with such "decisions from description", is seems hardly surprising that people fail to always decide optimally (Binmore, 1999; Plott, 1996).

A central idea in the preceding discussion was that preferences, rather than being immediately accessed, often need to be inferred. Various sources of information can support this inferential process, such as the decision context, emotions, and memory. These cues vary in their reliability and validity, and can sometimes conflict. In order to make a rational decision, the decision-maker will need to integrate them appropriately. Although cue integration can be difficult, it can be learned. Thus, if someone learns that current affective reactions are relatively poor predictors of future affective reactions, they may learn to give relatively more weight to other considerations (Beer, Knight, & D'Esposito, 2006).

As people gain more experience with types of decision problems, they should learn to make more optimal decisions. Even the expectation that a gamble will be played repeatedly can increase the tendency to choose reward maximizing alternatives (Lopes, 1981; Wedell & Böckeholt, 1990). According to Friedman (1998, p. 941) 'every choice "anomaly" can be greatly diminished or entirely eliminated in appropriately structured learning environments' and evidence discussed in this section suggests that this may indeed be the case.

Probability matching and maximizing

A phenomenon that has puzzled researchers for a long time and that is closely related to learning in decision making is probability matching. Consider a simple game in which you repeatedly choose to look behind one of two doors. There will be a reward behind one of the doors, but you don't know which one. Unbeknownst to you, the reward is placed behind door A 60% of the time, and behind door B for the remaining 40% of trials. To maximize your overall reward, you should always look behind door A. However, a consistent finding in these probability learning tasks is that people usually fail to display this maximizing behavior. Instead, they often probability-match, looking behind door A roughly 60% of the time, and behind door B roughly 40% of the time. Although results also indicate "overshooting" (choosing A somewhat more frequently than its probability of being correct, e.g., Friedman & Massaro, 1998), the proportion of maximizing responses is usually much less than 100%. Various explanations have been offered for this finding. For instance, the reward sequences used in many studies are not exactly random, in which case it is optimal to not always choose the same option (Fiorina, 1971). Also, the reward magnitudes are usually quite small, so that the difference between maximizing and matching may not be notable (e.g., Vulkan, 2000), and people may expect the reward schedules to change over time (Peterson & Ulehla, 1965; Wolford, Newman, Miller, & Wig, 2004), in which case it is good to sometimes check the previously inferior options (this is known as the exploration-exploitation trade-off). Finally, most studies only report group averages, which can mask large individual differences in strategies (Shanks, Tunney, & McCarthy, 2002). All these explanations seem reasonable and Shanks, Tunney, and McCarthy (2002) showed that a large majority of participants can learn to maximize when they are provided with large financial incentives, meaningful and regular feedback (informing participants how much they could have earned), and make many repeated decisions.

Experience vs description

In recent years, researchers have directly compared decisions from experience to decisions from description. Keeping within the monetary gamble framework, Barron and Erev (2003) replicated a number of classic experiments, with the crucial difference that rather than describing the probabilities and rewards associated with each alternative, participants could experience the gambles in a two-armed bandit task, repeatedly choosing which machine to play and observing a randomly drawn reward from the appropriate distribution. Interestingly, Barron and Erev found a number of results that directly oppose the findings of description-based studies. For instance, they found a reversed common ratio effect, in which the riskier of two options became less attractive after multiplying the probability by a common ratio. They also found more risk seeking in the gain than loss domain, and an underweighting rather than overweighting of small probabilities. Thus, a number of key results underlying Prospect Theory have not been replicated in experience-based decision making.

Why do experience-based decisions show opposite patterns to description-based decisions? With respect to the underweighting of small probabilities, it might be the result of limited sampling (Fox & Hadar, 2006; Hertwig, Barron, Weber & Erev, 2004; Rakow, Demes, & Newell, 2008). Because the sampling distribution for the binomial distribution is highly skewed for small samples and low probabilities, rare events are often under-represented and may not occur at all[5]. However, even when description-based decision scenarios exactly match experienced probabilities, differences between experience- and description-based decisions can be found (Hau, Pleskac & Herwig, 2010; Ungemach, Chater & Stewart, 2009; but see Rakow et al., 2008). Underweighting of small probabilities may also result from memory limitations and recency effects, resulting in biased samples from memory, or an exploration-exploitation trade-off in situations where only feedback from chosen options is received. However, neither of these seems sufficient to fully explain the experience-description gap (Hertwig & Erev, 2009; Camilleri & Newell, 2011). Camilleri and Newell (2011) showed that repeated, consequential decisions may play a crucial factor in underweighting. While reliable underweighting was found when participants repeatedly chose which gamble to play and received feedback on the outcome (whether only of the gamble played, or also that of the foregone gamble), this was not the case for a sampling paradigm in which participants first chose a number of gambles to observe before making a final decision. While it is likely that multiple factors play a role in the "experience-description gap", the findings clearly show that results from description-based studies cannot always be generalized to experience-based decision making.

Discovering preferences

When people are given the opportunity to learn, they can not only learn about the probabilities of outcomes, but also about what those outcomes are like. Plott (1996) argues that preferences are usually not immediately obvious and need to be discovered. According to his "discovered preference hypothesis", decisions may initially be subject to a variety of biases, but after repeatedly making a decision and experiencing its outcome, people will converge to rational decisions consistent with EUT. In contrast to strong claims of preference construction, he argues for stable underlying preferences which can be accessed after sufficient learning. Some evidence for this has been found. For instance, Hoeffler and Ariely (1999) show how experience in a decision environment results in more stable preferences, especially after repeated decisions. A study by Amir and Levav (2008) shows that whether preferences stabilize may depend on the types of decisions encountered. Some decision situations, such as those with an asymmetrically dominated alternative, may lead people to use choice rules with little consideration of attribute trade-offs. Such shallow processing seems to prohibit the formation of stable preferences. When situations require an explicit consideration of attribute trade-offs, and with that a more thorough preference inference process, preferences do tend to stabilize.

Research has shown that experience can reduce or even eliminate preference reversals and other violations of EUT. In the study by List (2002) described earlier, professional card dealers did not show a reliable preference reversal between separate and joint evaluation, although non-professional card dealers did. According to Berg, Dickhhaut, and Rietz (2010), preference reversals between choice and pricing can be greatly reduced when participants are given an explicit incentive to reveal their true preference. In a meta-analysis of studies in the Lichtenstein and Slovic (1971) paradigm, they show that with such incentives, behaviour is consistent with stable underlying preferences, although expressions of these in different tasks are subject to independent sources of noise. Other work has found that the general willingness to pay/willingness to accept gap (Coursey, Hovis, & Schulze, 1987; Shogren, Shin, Hayes, & Kliebenstein, 1994), the related endowment effect (List, 2004; Plott & Zeiler, 2005), and loss aversion (Myagkov & Plott, 1997) are eliminated or at least substantially diminished in repeated market settings with meaningful feedback and incentives. Similar results were found for preference reversals between choice and pricing (Braga & Starmer, 2003; Cox & Grether, 1996). Christensen, Heckerling, Mackesyamiti, Bernstein, and Elstein (1995) found much smaller framing effects for medical experts compared to novices. In addition, repeated play of gambles with outcome feedback reduces violations of EUT in Allais type problems (van de Kuilen and Wakker, 2006) and leads people to approach maximizing expected value (Keren & Wagenaar, 1987; Barron & Erev, 2003). More recently, van de

Kuilen (2009) conducted a study in which he found that the best fitting probability weighting of Prospect Theory approached linearity after increased experience.

Taken together, these results indicate that while deviations from EUT are often observed when people make decisions in novel and usually hypothetical situations, after sufficient learning with meaningful feedback and incentives, they can learn to respond in close accordance to EUT. Although there may be limits to this – for instance, due to the consistently found difficulty in predicting future preferences – there seems to be some scope for EUT as both a normative and descriptive theory of decision making.

Conclusion

During its long history, the scientific study of decision making has progressed from a mainly normative theory to a mainly descriptive account. Common violations of the axioms of Expected Utility Theory (EUT) have led many to question its descriptive validity. However, recent results show that a complete abandonment of normative principles may be premature. Violations are often found when people are presented with decisions between monetary gambles, or other abstract descriptions of decision problems. Research on how people learn to make decisions from experience has shown that, given sufficient practice, people can and often do learn to make decisions in accordance with the principles of EUT. In novel situations, people may need to infer their preferences from various sources, such as the immediate decision context, and their memory and emotions. This inferential process may initially give rise to highly variable preferences and decisions. But repeated exposure to decisions and their consequences may provide people with enough evidence to make informed decisions, pushing them in the direction of rationality.

Future directions

- Can repeated experience allow all violations of EUT to be eliminated?
- To what extent (if any) can decisions be guided by unconscious influences?
- How should differences across individuals in decision making (e.g., some people are more risk-averse than others) be understood?
- What can decision making research contribute to the remediation of pathological decision behavior such as addictions?
- How might sequential sampling models of one-shot decisions be extended to repeated decisions that involve learning?
- How does the experience/description contrast apply in more realistic decision problems? In particular, what about contexts where decision makers integrate both forms of information?

Acknowledgements

Preparation of this chapter was supported by the Economic and Social Research Council (ESRC), grant RES-062-23-1511. We are grateful to David Lagnado and Daniel Reisberg for providing invaluable comments on earlier drafts of this chapter.

Endnotes

1. The St Petersburg game involves tossing a fair coin until heads comes up, after which the game ends and the player receives a payout of $\$2^n$, where *n* denotes the total number of throws. The expected payout is

$$\left(\frac{1}{2}\right)\times 2+\left(\frac{1}{4}\right)\times 4+\left(\frac{1}{8}\right)\times 8+\ldots=\sum_{n=1}^{\infty}\left(\frac{1}{2}\right)^{n}2^{n}=\infty\ ,$$

but people generally are are typically willing to pay only about $10 to play the game.

2. The first result implies that $u(\$1m) > .1 \times u(\$5m) + .89 \times u(\$1m)$, which can be rearranged to give $.11 \times u(\$1m) > .1 \times u(\$5m)$. Yet the second result implies the opposite, namely that $.11 \times u(\$1m) < .1 \times u(\$5m)$.

3. The first result implies that $1/3 \times u(\$100) > P(\text{black}) \times u(\$100)$, so that $1/3 > P(\text{black})$. However, the second result implies that $(1/3 + (2/3 - P(\text{black}))) \times u(\$100) < 2/3 \times \text{u}(\$100)$, which is equivalent to $1 - P(\text{black}) < 2/3$, so that $P(\text{black}) > 1/3$.

4. One could argue that, insofar as decision satisfaction is the ultimate goal, this finding calls into question the normative status of maximizing utility. But this is not necessarily the case. While the results indicate that the utility of outcomes may be affected by the decision context, as also proposed in regret theory (Loomes & Sugden, 1982), one can still argue that maximizing context specific utility is the normative strategy. Moreover, the results from these studies are correlational, and maximizers and satisficers may differ on a number of personality characteristics which affect their satisfaction.

5. For instance, with 10 draws from a binomial distribution with p=.1, the probability of observing no win is roughly $P(X=0) = .35$, hence a large proportion of people will experience less than the expected number of wins.

References

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'ecole americaine, *Econometrica, 21,* 503-546.

Amir, O., & Levav, J. (2008). Choice construction versus preference construction: The instability of preferences learned in context. *Journal of Marketing Research*, *45*, 145-158.

Anderson, M. C., & Spellman, B. A. (1995). On the status of inhibitory mechanisms in cognition: Memory retrieval as a model case. *Psychological Review*, *102*, 68 -100.

Audley, R. J. (1960). A stochastic model for individual choice behavior. *Psychological Review*, *67*, 1-15.

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215-233.

Beattie, J., & Baron, J. (1991). Investigating the effect of stimulus range on attribute weight. *Journal of Experimental Psychology: Human Perception and Performance*, *17*, 571-585.

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, *50*, 7-15.

Bechara, A., Tranel, D., Damasio, H., & Damasio, A. H. (1996). Failure to respond autonomically to anticipated future outcomes following damage to prefrontal cortex. *Cerebral Cortex*, *6*, 215-225.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. H. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*, 1293-1295.

Beer, J.S., Knight, R.T., & D'Esposito, M. (2006). Controlling the integration of emotion and cognition: The role of the frontal cortex in distinguishing helpful from hurtful emotional information. *Psychological Science*, *17*, 448-453.

Bentham, J. (1948). *An introduction to the principle of morals and legislation.* Oxford, UK: Blackwell. (Originally work published in 1789)

Berg, J. E., Dickhaut, J. W., & Rietz, T. A. (2010). Preference reversals: The impact of truth-revealing monetary incentives. *Games and Economic Behavior*, *68*, 443-468.

Berger, J., Meredith, M., & Wheeler, S.C. (2008). Contextual priming: Where people vote affects how they vote. *Proceedings of the National Academy of Sciences*, *105*, 8846- 8849.

Bergert, F.B., & Nosofsky, R.M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 999-1019.

Bernouilli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, *22*, 23-36. (Original work published in 1738)

Binmore, K. (1999). Why experiment in economics? *The Economic Journal,* 109, F16-F24.

Braga, J., & Starmer, C. (2005). Preference anomalies, preference elicitation, and the discovered preference hypothesis. *Environmental and Resource Economics*, *32*, 55-89.

Brendl, C.M., Markman, A.B., & Messner, C. (2003). Devaluation of goal-unrelated choice options. *Journal of Consumer Research*, *29*, 463-473.

Bröder, A. (2000). Assessing the empirical validity of the “take-the-best” heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 1332–1346.

Bröder, A., & Newell, B.R. (2008). Challenging some common beliefs: Empirical work within the adaptive toolbox metaphor. *Judgment and Decision Making*, *3*, 205-214.

Busemeyer, J.R., & Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*, 432-459.

Camerer, C.F., & Ho, T.H. (1994). Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty*, *8*, 167-196.

Camilleri, A.R. & Newell, B.R. (2011). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review, 18,* 377-384.

Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (2003). Fast, frugal and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes*, *90*, 63-86.

Christensen, C., Heckerling, P., Mackesyamiti, M.E., Bernstein, L.M., & Elstein, A.S. (1995). Pervasiveness of framing effects among physicians and medical students. *Journal of Behavioral Decision Making*, *8,* 169-180.

Coursey, D.L., Hovis, J.L., & Schulze, W.D. (1987). The disparity between willingness to accept and willingness to pay measures of value. *Quarterly Journal of Economics*, *102*, 679-690.

Cox, J.C., & Grether, D.M. (1996). The preference reversal phenomenon: Response mode, markets and incentives. *Economic Theory*, *7*, 381-405.

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

Damasio, A.R. (1996). The somatic marker hypothesis and the possible functions of the prefrontral cortex. *Philosophical Transactions of the Royal Society of London (Series B), 351,* 1413-1420.

Damasio, H., Grabowski, T., Frank, R., Galaburda, A.M., & Damasio, A.R. (1994). The return of Phineas Gage: Clues about the brain from the skull of a famous patient. *Science*, *264*, 1102-1105.

Dunn, B.D., Dalgleish, T., & Lawrence, A.D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience and Biobehavioral Reviews*, *30*, 239-271.

Edwards, K. D. (1996). Prospect theory: A literature review. *International Review of Financial Analysis*, *5*, 19-38.

Ellsberg, D. (1961). Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics*, *75*, 643-669.

Finucane, M. L., Alhakami, A., Slovic, P., & Johnson, S. M. (2000). The affect heuristic in judgments of risks and benefits. Journal of Behavioral Decision Making, 13, 1–17.

Finucane, M. L., Peters, E. & Slovic, P. (2003). Judgment and decision making: The dance of affect and reason. In S. L. Schnieider & J. Shanteau (Eds.), *Emerging perspectives on judgment and decision research* (pp. 327–364). Cambridge, UK: Cambridge University Press.

Fiorina, M.P. (1971). A note on probability matching and rational choice. *Behavioral Scienc*e, *16,* 158-166.

Fox, C. R., & Hadar, L. (2006). "Decisions from experience" = sampling theory + prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, *1*, 159–161.

Frederickson, B.L. & Kahneman, D. (1993). Duration neglect in retrospective evaluations of affective responses. *Journal of Personality and Social Psychology*, *65*, 45-55.

Friedman, D. (1998). Monty Hall's three doors: construction and deconstruction of a choice anomaly. *American Economic Review*, *88,* 933-946.

Friedman, D., & Massaro, D.W. (1998). Understanding variability in binary and continuous choice. *Psychonomic Bulletin & Review*, *5,* 370-389.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451-482.

Gigerenzer, G. & Goldstein, D.G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650-669.

Gigerenzer, G., Todd, P.M., & the ABC Research Group (1999). *Simple heuristics that make us smart*. New York, NY: Oxford University Press.

Gilbert, D.T., Gill, M., & Wilson, T.D. (1998). The future is now: Temporal correction for affective forecasting. *Organizational Behavior and Human Decision Processes*, *88*, 430-444.

Goldstein, D.G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, *109*, 75-90.

Hau, R., Pleskac, T. J., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making*, *68*, 48-68.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*, 534-539.

Hertwig, R. & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Science*, *13*, 517-523.

Hoeffler, S., & Ariely, D. (1999). Constructing stable preferences: A look into dimensions of experience and their impact on preference stability. *Journal of Consumer Psychology*, *8*, 113-139.

Hoyer, W.D. & Brown, S.P. (1990) Effects of brand awareness on choice for a common, repeat-purchase product. *Journal of Consumer Research*, *17*, 141-148.

Hsee, C. K. (1996). The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives. *Organizational Behavior and Human Decision Processes*, *67*, 247-257.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*, 90-98.

Iyengar, S.S., Wells, R.E., & Schwarz, B. (2006). Doing better but feeling worse: Looking for the "best" job undermines satisfaction. *Psychological Science*, *17*, 143-150.

Johnson E.J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: a query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, And Cognition*, *33*, 461-474.

Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.). *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press. p. 49-81.

Kahneman, D., & Snell, J. (1992). Predicting a changing taste: Do people know what they will like? *Journal of Behavioral Decision Making*, *5*, 187-200.

Kahneman, D., Knetsch, J.L., & Thaler, R.H. (1990). Experimental tests of the endowment effect and the Coase theorem. *The Journal of Political Economy*, *98*, 1325-1348.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263-292.

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

Kahneman, D., & Tversky, A. (2000). *Choices, values, and frames*. New York, NY: Cambridge University Press.

Kahneman, D., Wakker, P. P., & Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. *The Quarterly Journal of Economics*, *112*, 375–405.

Keeny, R.L., & Raiffa, H. (1976). Decisions with multiple objectives: Preferences and value tradeoffs. New York, NY: John Wiley & Sons.

Keren, G., & Schul, Y. (2008). Two is not always better than one: A critical evaluation of two-system theories. *Perspectives on Psychological Science*, *4*, 533-550.

Keren, G., & Wagenaar, W.A. (1987). Violation of utility theory in unique and repeated gambles. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 387-391.

Knetsch, J.L. (1989). The endowment effect and evidence of nonreversible indifference curves. *The American Economic Review*, *79*, 1277-1284.

Krantz, D.H., & Kunreuther, H.C. (2007). Goals and plans in decision making. *Judgment and Decision Making*, *2*, 137-168.

Krantz, D.H., Luce, R.D., Suppes, P., & Tversky, A. (1971). *Foundations of measurement (Volume 1): Additive and polynomial representations*. New York, NY: Academic Press.

Lichtenstein, S., & Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology*, *89*, 46-55.

List, J.A. (2002). Preference reversals of a different kind: The ‘more is less’ phenomenon. *American Economic Review*, *92*, 1636-1643.

List, J.A. (2004). Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica*, *72*, 615-625.

Loewenstein, G. & Adler, D. (1995). A bias in the prediction of tastes. *The Economic Journal*, *105*, 929-937.

Loewenstein, G, O'Donoghue, T., & Rabin, M. (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics*, *118*, 1209-1248.

Loewenstein, G., Weber, E. U., Hsee, C. K., & Welch, E. S. (2001). Risk as feelings. *Psychological Bulletin*, *127*, 267–286.

Loomes, G., & Sugden, R. (1982). Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal, 92*, 805-824.

Lopes, L.L. (1981). Decision making in the short run. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 377-385.

Maia, T. V., & McClelland, J. L. (2004). A re-examination of the evidence for the somatic marker hypothesis: What participants know in the Iowa gambling task. *Proceedings of the National Academy of Sciences*, *101*, 16075-16080.

March, J.G. (1978). Bounded rationality, ambiguity, and the engineering of choice. *The Bell Journal of Economics*, *9*, 587-608.

Markman, A.B., Brendl, C.M., & Kim, K. (2007). Preference and the specificity of goals. *Emotion*, *7*, 680-684.

Markman, A. B., & Medin, D. L. (2002). Decision making. In D.L. Medin & H. Pashler (Eds.), Stevens Handbook of Experimental Psychology, Volume 2 (3rd ed., pp. 413-466).

Mellers, B.A., Cooke, A.D.J. (1994). Trade-offs depend on attribute range. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 1055-1067.

Myagkov, M., & Plott, C.R. (1997). Exchange economies and loss exposure: Experiments exploring prospect theory and competitive equilibria in market environments. *American Economic Review*, *87*, 801-828

Neumann, J., von & Morgenstern, O. (1947). *Theory of Games and Economic Behavior* (2nd edition). Princeton, NJ. Princeton University Press.

Newell, B. R., Lagnado, D. A., & Shanks, D. R. (2007). *Straight choices: The psychology of decision making*. Hove, UK: Psychology Press.

Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 53-65.

Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone ''takes-the-best''. *Organizational Behavior and Human Decision Processes*, *91*, 82–96.

North, A. C., Hargreaves, D. J., & McKendrick, J. (1997). In-store music affects product choice. *Nature*, *390*, 132.

Oeusoonthornwattana, O., & Shanks, D. R. (2010). I like what I know: Is recognition a non-compensatory determiner of consumer choice? *Judgment and Decision Making*, *5*, 310-325.

Ortmann, A., Gigerenzer, G., Borges, B., & Goldstein, D.G. (2008). The recognition heuristic: A fast and frugal way to investment choice? In C.R. Plott & V.L. Smith (Eds.) *Handbook of Experimental Economics Results: Vol. 1* (Handbooks in Economics No. 28), pp. 993–1003. Amsterdam: North-Holland.

Pachur, T., Bröder, A., & Marewski, J.N. (2008). The recognition heuristic in memory-based inference: Is recognition a non-compensatory cue? *Journal of Behavioral Decision Making*, *21*, 183-210.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology*, *43*, 87-131.

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker.* Cambridge, UK: Cambridge University Press.

Peterson, C. R., & Ulehla, Z. J. (1965). Sequential patterns and maximizing. *Journal of Experimental Psychology*, *69*, 1-4.

Plott, C.R. (1996). Rational individual behavior in markets and social choice processes: The discovered preference hypothesis, in Arrow, K., Colombatto, E., Perleman, M. & Schmidt, C. (Eds.) *Rational Foundations of Economic Behavior*. London: Macmillan and St. Martins. 225–250.

Plott, C.R., & Zeiler, K. (2005). The willingness to pay-willingness to accept gap, the endowment effect, subject misconceptions, and experimental procedures for eliciting valuations. *American Economic Review*, *95*, 530-545.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization, 3*, 323-343.

Rakow, T., Demes, K.A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes*, *106*, 168–179.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59-108.

Redelmeier, D. A., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, *66*, 3-8.

Rieskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics and how can we tell? In G. Gigerenzer, P. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 141-168). Oxford, UK: Oxford University Press.

Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207-236.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, *108*, 370-392.

Rottenstreich, Y., & Hsee, C. K. (2001). Money, kisses and electric shocks: On the affective psychology of probability weighting. *Psychological Science*, *12*, 185–190.

Samuelson, P. A. (1963). Risk and uncertainty: A fallacy of large numbers. *Scientia*, *98*, 108-113.

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making, *Journal of Risk and Uncertainty*, *1*, 7-59.

Savage, L. J. (1954). *The foundations of statistics.* New York, NY: Dover.

Schoemaker, P. (1982). The expected utility model: Its variants, purposes, evidence and limitations. *Journal of Economic Literature*, *20*, 529-563.

Schwartz, B., Ward, A., Monterosso, J., Lyubomirsky, S., White, K., & Lehman, D.R. (2002). Maximizing versus satisficing: Happiness is a matter of choice. *Journal of Personality and Social Psychology*, *83*, 1178–1197.

Sen, A. (1993). Internal consistency of choice. *Econometrica*, *61*, 495–521.

Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition*, *49*, 11-36.

Shanks, D. R., Tunney, R. J. & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, *15*, 233-250.

Shefrin, H. M., & Thaler, R. H. (1992). Mental accounting, saving, and self- control. In G. Loewenstein & J. Elster (Eds.), *Choice over time* (pp. 287-330). New York: Russell Sage Foundation.

Shogren, J.F., Shin, Seung, Y., Hayes, Dermot, J. and Kliebenstein, J., (1994) Resolving differences in willingness to pay and willingness to accept. *American Economic Review, 84*, 255–270.

Simon, H.A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*, 99-118.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63,* 129-138.

Simonson, I. (1989). Choice based on reasons: The case of attraction of compromise effects. *Journal of Consumer Research*, *16*, 158-174.

Simonson, I. (1990). The effect of purchase quantity and timing on variety seeking behavior. *Journal of Marketing Research*, *27*, 150-162.

Simonson, I. (2008). Will I like a “medium” pillow? Another look at constructed and inherent preferences. *Journal of Consumer Psychology*, *18*, 155-169.

Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, *119*, 3–22.

Slovic, P. (1975). Choice between equally valuable alternatives. *Journal of Experimental Psychology: Human Perception and Performance*, *1*, 280-287.

Slovic, P. (1995). The construction of preference. *American Psychologist*, *50*, 364-371.

Smith, P. L. (2000). Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, *44*, 408-463.

Steward, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1-26.

Thaler, R. (1980). Towards a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, *1*, 39-60.

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science*, *4*, 199-214.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*, 281-299.

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

Tversky, A. & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, *59*, 251-278.

Tversky, A. & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *323*, 297-323.

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, *39,* 1179-1189.

Tversky, A., Stattath, S. & Slovic, P. (1988). Contingent weighting in judgment and choice. *Psychological Review*, *95*, 371-384.

Tversky, A., & Thaler, R. H. (1990). Anomalies: Preference reversals. *Journal of Economic Perspectives*, *4*, 201-211.

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, *20*, 473-479.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108, 550-592.

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, *111*, 757-769.

van de Kuilen, G. (2009). Subjective probability weighting and the discovered preference hypothesis. *Theory and Decision*, *67*, 1-22.

van de Kuilen, G. & Wakker, P.P. (2006). Learning in the Allais paradox. *Journal of Risk and Uncertainty*, *33*, 155-164.

Vulkan, N. (2000). An economist's perspective on probability matching. *Journal of Economic Surveys*, *14*, 101-118.

Wakker, P. P. (1989). Additive representations of preferences: a new foundation of decision analysis. Dordrecht, NL: Kluwer Academic.

Weber, E. U., & Johnson, E. J. (2006). Constructing preferences from memories. In S. Lichtenstein & P. Slovic (Eds.), *The construction of preference* (pp. 397–410). New York: Cambridge University Press.

Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, *60*, 53-85.

Wedell, D. H., & Böckenholt, U. (1990). Moderation of preference reversals in the long run. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 429-438.

Wilson, T.D. & Gilbert, D.T. (2005). Affective forecasting: Knowing what to want. *Current Directions in Psychological Science*, *14*, 131-134.

Wirtz, D., Kruger, J., Scollon, C.N., & Diener, E. (2003). What to do on spring break? The role of predicted, on-line, and remembered experience in future choice. *Psychological Science*, *14*, 520-524.

Wolford, G., Newman, S. E., Miller, M. B., & Wig, G. S. (2004). Searching for patterns in random sequences. *Canadian Journal of Experimental Psychology*, *58*, 221-228.

Wu, G. & Gonzalez, R. (1996). Curvature of the probability weighting function. *Management Science*, *42*, 1676-1690

Figure captions

Figure 1: The value function in Prospect Theory.

Figure 2: The weighting function (*w*) for cumulative probabilities (*p*) in Cumulative Prospect Theory.

Figure 3: Schematic representation of the decision alternatives in research on context effects. The direction of preference is upwards and to the right.

Figure 4: Hypothetical evidence accumulation process for three options.

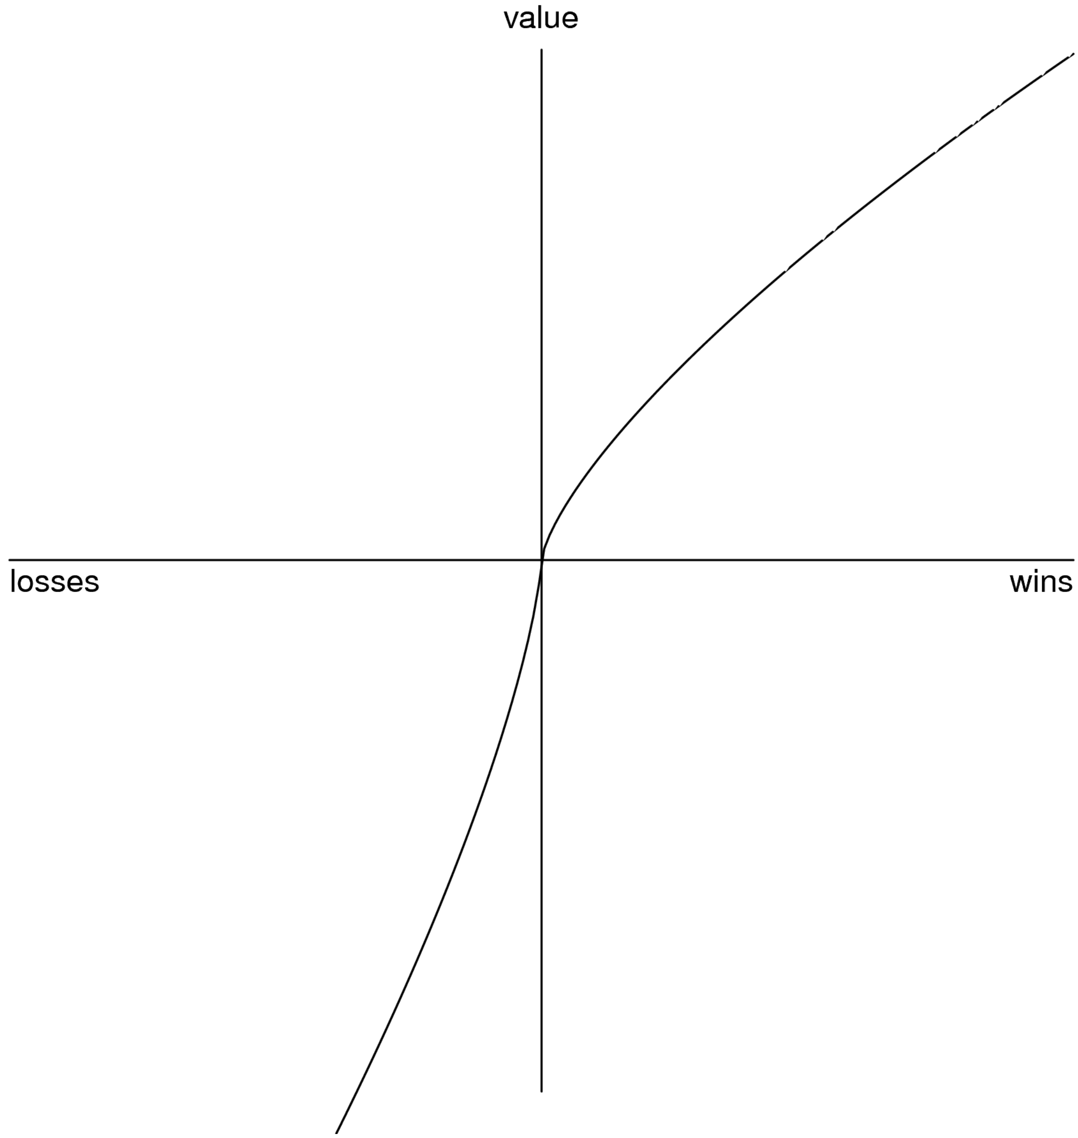

Figure 1

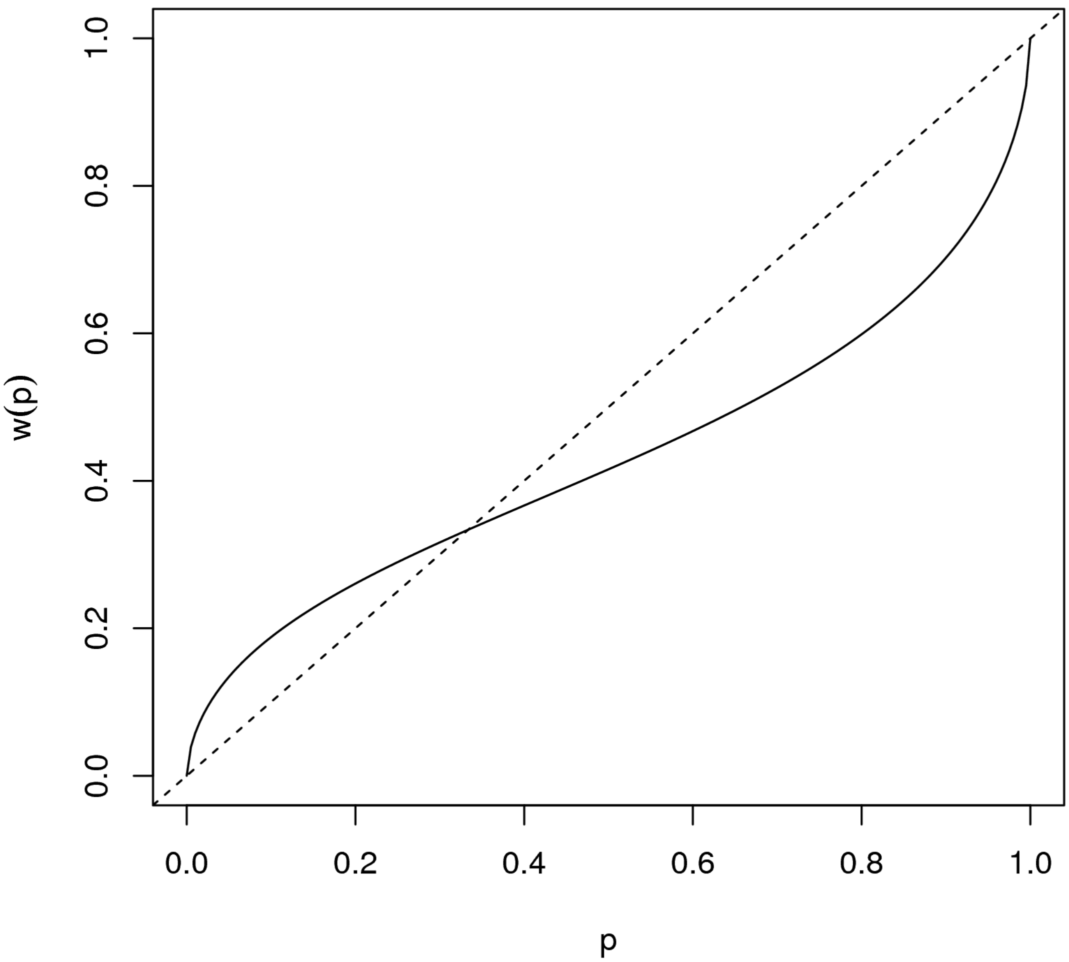

Figure 2

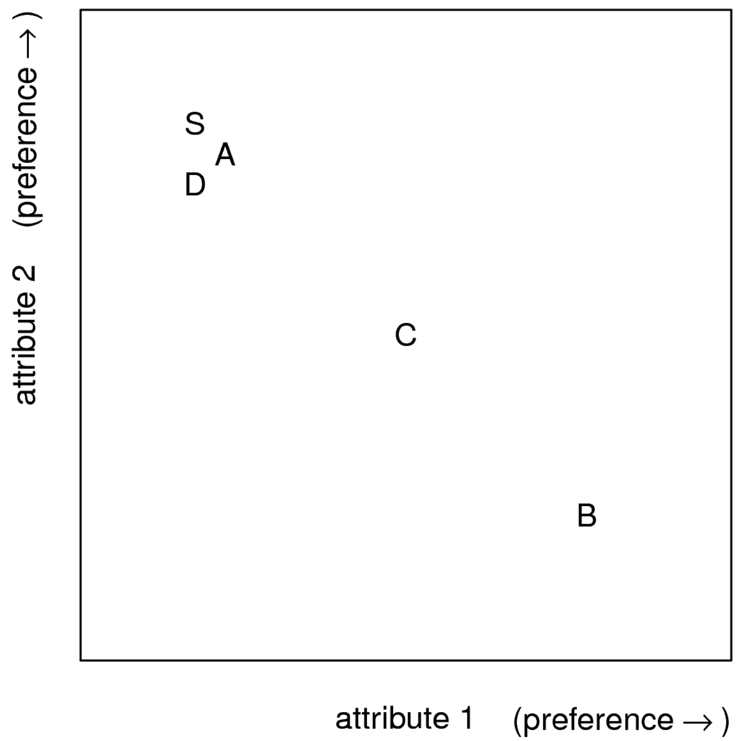

Figure 3

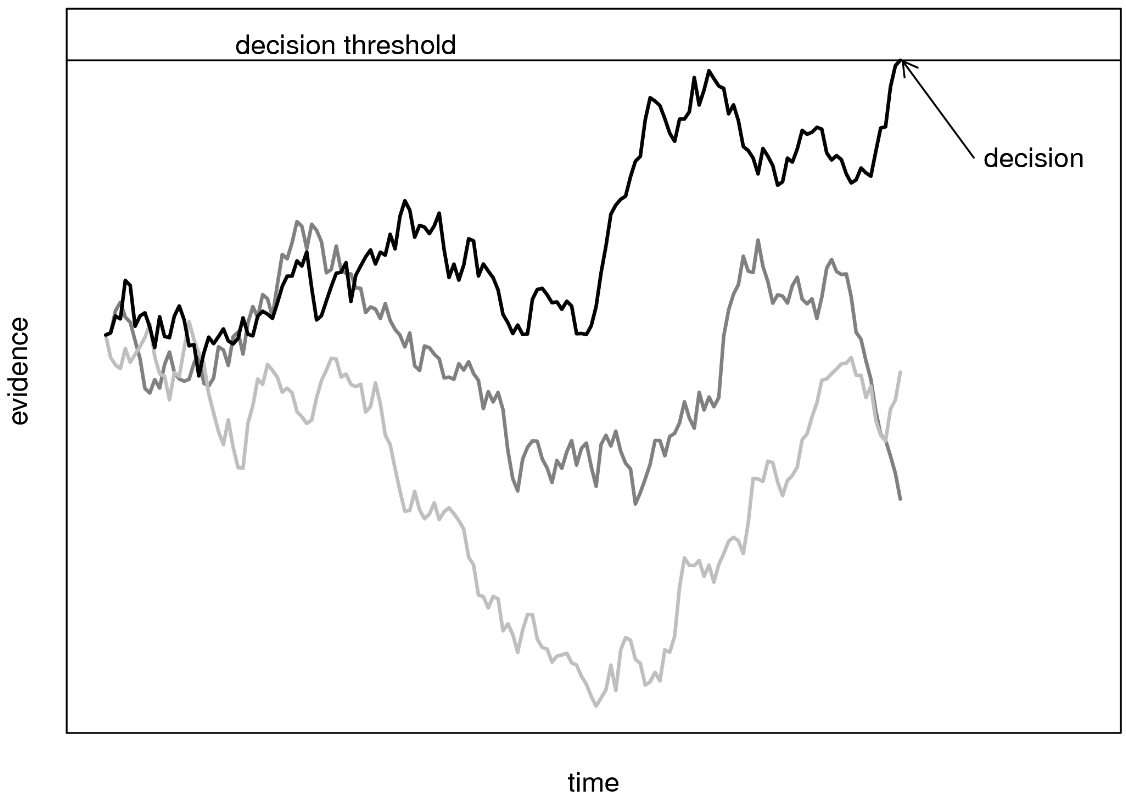

Figure 4