Biases in thinking and decision-making

Inquiry questions

- Do people tend to take shortcuts when thinking and making decisions?
- What kind of simplified decisions strategies do they use?
- Do these strategies lead to biased or irrational decisions?
- Are these biases predictable?

What you will learn in this section

- System 1 and system 2 thinking
  - Heuristics are cognitive shortcuts or simplified strategies; heuristics lead to cognitive biases which may be identified by comparing the decision to a normative model
  - System 1: immediate automatic responses; system 2: rational deliberate thinking (Daniel Kahneman)
- Common causes of intuitive thinking
- The tendency to focus on a limited amount of available information
  - This tendency is the result of selective attention and schemas
  - Asymmetric dominance (Huber, Payne and Puto, 1982)
    - Explanations of asymmetric dominance: selective attention and a desire to justify one's choice
  - Framing effect (Tversky and Kahneman, 1981)
    - Prospect theory: individuals think about utilities as changes from a reference point
    - Violation of the normative expected-utility theory
    - "Avoid risks, but take risks to avoid losses"
- The tendency to seek out information that confirms pre-existing beliefs
  - Confirmation bias. Wason's four-card problem (Wason, 1968)
  - Congruence bias (Tschirgi, 1980; Wason, 1960): people are trying to obtain positive results rather than useful information
  - Illusory correlation and implicit personality theories (Chapman and Chapman, 1969): illusory correlations based on prior beliefs are stable and resistant to change
- The tendency to avoid the mental stress of holding inconsistent cognitions
  - The theory of cognitive dissonance
    - Belief disconfirmation paradigm (Festinger, 1956)
    - Induced-compliance paradigm (Freedman and Fraser, 1966)

This section also links to:

- principles of the cognitive approach to behaviour—cognitive processes do not function in isolation; biases in cognitive processes can be systematic and predictable
- thinking and decision-making (normative and descriptive models; limited computational capacity, the influence of emotion on thinking, meta-goals of decision-making)
- stereotypes (sociocultural approach to behaviour).

Psychology in real life

An interesting case of friendly competition between humans and machines is the Supreme Court Forecasting Project 2002: http://wusct.wustl.edu/

It compared the accuracy of two different ways to predict outcomes of the Supreme Court cases in the USA: informed opinions of 83 legal experts versus a computer algorithm. They were predicting in advance the votes of each of the nine individual justices for every case in the Supreme Court in 2002. The same algorithm was used to predict outcomes of all cases, but legal experts were only predicting the cases that were within their area of expertise.

The computer algorithm seemed very reductionist. It only took into account six simple factors such as the issue area of the case or whether or not the petitioner argued that a law or practice is unconstitutional.

Both predictions were posted publicly on a website prior to the announcement of each of the Court's decisions. There was a lot of suspense.

The experts lost the game: the computer correctly predicted 75% of the Supreme Court decisions, while the experts collectively made only 59.1% of correct predictions. Note that all the decisions were binary (affirm/reverse), so the experts did only 9.1% better than what could be achieved by a toss of a coin.

Why can human experts who have access to detailed information about a case turn out to be such bad predictors? What is it about human decision-making that allows a simple computer algorithm to outperform collective wisdom of people with years of education and experience behind them? If human decisions are biased, can these biases be fixed?

System 1 and system 2 thinking

We have already discussed the important distinction between **normative models** (for example, logic, theory of probability, utility theory) and **descriptive models** of thinking and decision-making. Attractive as they are in leading us to the most rational choice, normative decision theories are unrealistic when it comes to making decisions in real life. As already discussed, they do not account for:

- limited computational capacity
- the influence of emotion on thinking
- other goals that the decision-maker might have, for example, justifying the choice to others, confirming one's own belief or supporting self-esteem.

So, naturally, people use shortcuts and incomplete, simplified strategies which are known as **heuristics**. Heuristics may also be expressed as rules, which makes them an exciting area of research. If we identify and describe a set of common heuristics and prove that people actually use them in real-life decision-making scenarios, we will be able to predict what people are likely to think or do in certain situations. Moreover, we might be able to design computer intelligence that mimics human intelligence.

TOK

1. According to some philosophers, science has four functions: to describe, to explain, to predict and to control.

   To what extent can we describe, explain, predict and control human decisions?
2. Do you believe that artificial intelligence can be developed to the level where it can mimic human decisions, that is, predict what decisions humans will make in a particular situation?

Using heuristics leads to **cognitive biases** (which can be described effectively if you compare heuristics to the normative model for a particular situation). However, heuristics are useful. First, they save energy; we don't have to meticulously analyse all the aspects of the situation every time we are faced with a choice. Second, heuristics are often based on experience, which means that you used them before and it worked reasonably well. Of course the rule saying "if it worked before, it will work now" is not perfect, but it is reasonable enough for a variety of everyday situations.

Daniel Kahneman in 2003 proposed an extension to the information-processing approach by differentiating between **two independent systems**, system 1 and system 2.

This differentiation has become the core of his bestselling book *Thinking, Fast and Slow* (2011), which is a must-read if you have an interest in cognitive biases and behavioural economics. According to the theory, **system 1** thinking is fast, instinctive, emotional, automatic and relatively unconscious, whereas **system 2** thinking is slower, more analytical, logical, rule-based and conscious. System 1 is commonly referred to as "intuition".

Exercise

Go online and search for "behavioural economics". What is it? What are the main landmarks of research in this area?

Present your findings briefly in class.

It has been argued that system 1 developed as an adaptive reasoning mechanism which is based on prior experience (and survival goals) and enables us to make fast and reasonably accurate decisions that have proved to be sufficiently successful in the past. System 2 evolved later with the development of language and abstract reasoning, and this enables us to overcome some of our immediate automatic responses and analyse the situation in greater depth.

Due to this legacy, we use system 1 in the majority of common situations, but we switch to system 2 when the situation is unusual and complex or when we encounter difficulties with our intuitive response. By this reasoning, our thinking works **sequentially**: first, there is a fast and automatic system 1 response, and then this response is (or is not) corrected by the more conscious cognitive efforts of system 2.

System 1 works better in "predictable" environments. Arguably, in today's world with its high degree of complexity, tremendous rates of the production of new knowledge and rapidly changing circumstances, individuals need to be much more flexible and adapt more quickly. So the cognitive demands placed on system 2 processing seem to be increasing. This makes the study of heuristics and biases associated with system 1, as well as the way in which descriptive models of thinking deviate from normative models, even more pertinent.

Common causes of intuitive thinking

There have been numerous attempts to identify the common heuristics and cognitive biases, starting in the 1970s and, as the area gained more and more popularity, continuing in the three decades that followed. This has given rise to attempts at classification and a search for common causes of intuitive (automatic) thinking.

There is no universally accepted classification of heuristics or the common causes underlying them. However, we will focus on several major factors as examples to take a peek into the sheer variety:

- the tendency to focus on a limited amount of available information (asymmetric dominance, framing effect)
- the tendency to seek out information that confirms pre-existing beliefs (confirmation bias, congruence bias, illusory correlations and implicit personality theories)
- the tendency to avoid the mental stress of holding inconsistent cognitions (cognitive dissonance).

The tendency to focus on a limited amount of available information

ATL skills: Self-management

Before you read on, remember four meta-goals in the adaptive decision-maker framework discussed earlier in this chapter. If you find it difficult, go back and review the section "Thinking and decision-making".

As you continue reading, think about which of the four meta-goals can explain the tendency to focus on a limited amount of available information. You will find an answer later in the text.

Remember what you know about sensory memory and how it gets transferred into short-term memory (STM). Sensory memory has a high capacity, but a very limited duration. In order for information to reach STM, this information has to be attended to. However, we cannot pay attention to many chunks of information at the same time. From the sea of stimuli around us we have to single out one stimulus that we focus on and process further. This is known as **selective attention**.

Remember what you know about schemas (top-down processing), and you will realize that selective attention is mediated by our existing preconceptions and expectations—while travelling from sensory memory to STM, information passes through a "lens" of schemas. This may distort information or, at the very least, filter out certain aspects or details, however important they may be from the point of view of a normative model.

Let's look at several examples of cognitive biases that are rooted in this inherent tendency to focus on a limited amount of available information: asymmetric dominance and framing effect.

Asymmetric dominance

Huber, Payne and Puto (1982) studied decisions involving an **asymmetrically dominated decoy**. An example will make it clear what this is.

Suppose you are choosing between two alternatives, A and B, and you are basing your decision on two dimensions. For example, you are choosing between two holiday destinations (A = Swiss Alps and B = Goa), and there are two parameters that are important to you, affordability (you prefer cheaper) and ski opportunities (you like skiing more than swimming in the sea). In Graph A on the right, neither the target nor the competitor dominate the other, because one is better on one dimension and the other is better on the other dimension. You would prefer Goa because it is cheaper, but at the same time you would prefer the Swiss Alps because there is a skiing opportunity. Now what if we introduce the third option, a decoy? Imagine your third choice is C = Seychelles (see Graph B). For the sake of this example, let's say that having a holiday in the Seychelles is slightly more affordable than in the Swiss Alps, at the same time skiing is not a possibility, just as in Goa. This alternative (the Seychelles) is said to be an asymmetrically dominated decoy, because it is dominated by Goa (the same skiing opportunities, however more affordable), but not dominated by Swiss Alps (better skiing opportunities, but less affordable).

Note how the decoy is not a desirable choice. Of the three options, why would you choose it? If you attach more importance to skiing opportunities, you would choose the Alps, if you attach more importance to affordability, you would go for Goa. So why is it important?

It turns out that adding an asymmetrically dominated decoy in the choice set increases the probability of choosing the option that dominates the decoy. In our example, it means that people who are given options A, B and C would be more likely to choose B, compared with people who are given only two options, A and B. Note that nobody chooses C, but the mere presence of C in the choice set increases the probability of the choice of B.

Similarly, if we wanted more people to choose the Swiss Alps, we might have manipulated the choice set differently. In the choice set A, B, D (see Graph C below) D is the northern part of the Ural Mountains in Central Russia. For the sake of this example, let's say this destination is as expensive as the Swiss Alps (given the remoteness of the place), and there are plenty of skiing opportunities, but the skiing facilities are not as modern as in the Alps. In this new choice set, D is the asymmetrically dominated decoy (which people will not choose). It is dominated by the Swiss Alps (better facilities, the same in terms of affordability) but it is not dominated by Goa (better affordability, but less skiing opportunities). So, the theory predicts that in this choice set people would be more likely to prefer the Alps.

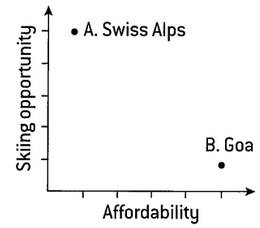

▲ Figure 3.17 Asymmetrically dominated decoys

Huber, Payne and Puto confirmed this prediction in a sample of 153 students who were required to make choices in six categories: cars, restaurants, beers, lotteries, films and television sets. Just as in our examples, the decision environments included two or three alternatives, with each alternative defined on two attributes. Overall, the predicted choice reversals were not large (3–9% of participants **switched their choice** in the predicted direction when the decoy was added to the choice set), but statistically significant.

Note that asymmetric dominance violates the normative model of rational choice, as there is no logical reason why the decoy can change your preferences. How can this result be explained?

- First, there's an explanation in terms of selective attention. We **selectively attend to one of the attributes** (for example, affordability) and ignore the other. We seem to be incapable of processing both the attributes simultaneously and look for an "excuse" to simplify the problem and process only one.
- Second, one might explain the observed data by the desire of the decision-maker to justify his or her choice. Arguably, the decoy gives you one more reason to say, "I have chosen X, because ...".

At this point you might remember our discussion of the adaptive decision-maker framework in which maximizing the ease of justification of a decision is one of the four meta-goals of decision-making.

A difference of 3–9% may become hugely important in large corporations looking for a competitive edge. Huber, Payne and Puto give the following example of a practical application of asymmetric dominance effect. A store owner has two camel hair jackets priced at $100 and $150 and finds that the more expensive jacket is not selling. A new camel hair jacket is added and displayed for $250; the new jacket does not sell, but sales of the $150 jacket increase. Think about more practical implications, for example, the choice of burgers at your local McDonalds. Potentially using the asymmetric dominance effect in marketing can bring millions in profit, and it does.

ATL skills: Social

Get into small groups and think of other practical applications of this cognitive bias. Use your knowledge from other subjects, such as history or business management.

Framing effect

The most influential normative model of choice under uncertainty is **expected utility theory**. In this theory you multiply the utility of an outcome by the probability of that outcome, and choose the outcome that yields the highest number. For example, suppose you were choosing between two gambles: if you choose option A, you get $10 for certain; if you choose option B, you get $200 with 6% probability. According to the normative theory, it is more rational to take a risk: the expected utility of option A is 10 * 1 = $10, whereas the expected utility of option B is 200 * 0.06 = $12.

However, numerous studies have demonstrated that in their real-life choices people do not always adhere to the predictions of the normative model. They seem to be too eager to take risk in some situations and too avoidant of risks in others, depending on seemingly irrelevant factors.

In 1979 Daniel Kahneman and Amos Tversky proposed a descriptive theory of decision-making under risk that is known as **prospect theory**. The idea behind the theory was to take the normative expected utility model and to modify it as little as possible to explain the observed deviations from the normative model. They were successful, and prospect theory quickly gained popularity as a descriptive model of choice.

Prospect theory claims that **individuals think about utilities as changes from a reference point** (and the reference point may be easily changed by the way the problem is formulated).

In one of their famous experiments, **Tversky and Kahneman (1981)** gave their subjects the following problem.

> Imagine that the USA is preparing for an outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimate of the consequences of the program are as follows:
>
> (and the options were different for two independent groups of subjects)

Group 1	Group 2:
Program A: 200 people will be saved.	**Program A:** 400 people will die.
Program B: there is 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved.	**Program B:** there is 1/3 probability that nobody will die, and 2/3 probability that 600 people will die.

▲ Table 3.9 Response options for the two groups in Tversky and Kahneman (1981)

Note that both choice sets are identical. The only difference is in how the situation is described, either in terms of **potential gains** ("will be saved") or in terms of **potential losses** ("will die"). Having said that, it is interesting that participants' choices in these two groups were reversed. Here is the percentage of individuals who chose each of the two programs in the two groups:

	Group 1	Group 2
Program A	72%	22%
Program B	28%	78%

▲ Table 3.10 Findings from Tversky and Kahneman (1981)

Since nothing changed in the problem from the rational point of view, this reversal of choices cannot be explained by the normative (expected utility) model. The expected utilities of the two programs are the same. So how can we explain the deviation from the normative model?

Presumably most individuals in group 1 choose program A because they are trying to avoid the risk of not saving anyone at all (2/3 probability). Whereas in group 2 individuals seemed to be more willing to take the risk: 400 deaths seems almost as bad as 600 deaths, so why not take a chance?

Tversky and Kahneman explain this finding in terms of a **shift in the reference point**. In the first version the reference point is the future state (600 people dead), so the options are perceived as potential **gains** (how many people can I save?). In the second version the reference is shifted to the present state (no one has died yet), so the options are perceived as potential **losses** (how many people can we lose?).

So, depending on whether outcomes are described ("framed") as gains or losses, subjects give different judgments: they are more willing to take risks to avoid losses and have a tendency to avoid risk associated with gains. In other words, "Avoid risks, but take risks to avoid losses" (Baron, 2008, p 270). This is known as **the framing effect**.

The prospect theory proposes a mathematical way to describe such deviations from the normative model. Schematically it is shown in Figure 3.18. In terms of the normative model, it does not matter where the reference point is: the value function should be a straight line. However, people assign less (positive) value to gains and more (negative) value to losses. We selectively redistribute our attention to the potential outcomes based on how the problem is framed.

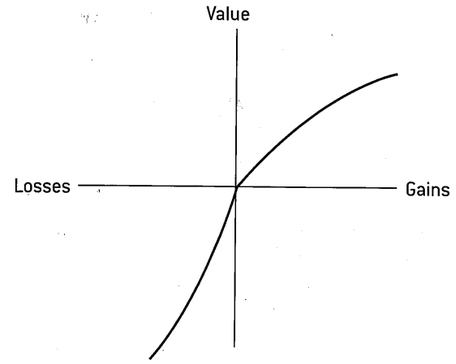

▲ Figure 3.18 Value function in prospect theory

ATL skills: Communication

How can using the framing effect be translated into monetary gains?

Think about how insurance agents frame their offers ("If you buy this insurance package you may avoid losing ...").

Can you identify other examples (commercials, business offers, job vacancies)? Think of some and share with the group.

The tendency to seek out information that confirms pre-existing beliefs

Confirmation bias

TOK

What do you know about confirmation bias from TOK classes? How does it create an obstacle in the production of knowledge?

Another common source of heuristics is the tendency to seek out information that confirms pre-existing beliefs. You might have discussed the role of confirmation bias in your theory of knowledge classes. Let's look at the psychological research behind this phenomenon.

A good illustration of confirmation bias as a logical heuristic (that is, a heuristic that violates logic as the normative model) is **Wason's four-card problem (1968)**.

Suppose you have the following four cards and you know that each card has a letter on one side and a number on the other side. You are also given the following rule: "If a card has a vowel on one side, then it has an even number on the other side". Your task is to "name those cards, and only those cards, that need to be turned over in order to determine whether the rule is true or false".

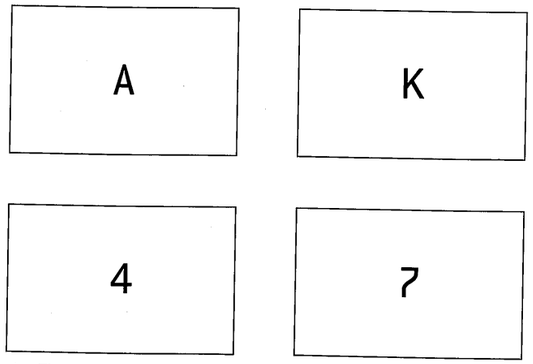

▲ Figure 3.19 Wason's four-card problem

Which cards will you turn over? And why?

The most popular answers in Wason's study were either "A only" or "A and 4". To see what the correct answer is, follow this logic.

- If you turn over the first card and there is an even number on the other side, it will support the rule. If there's an odd number on the other side, it will refute the rule.
- If you turn over the second card and there is an even number on the other side, it tells you nothing about the rule. Neither does an odd number. Since the letter is a consonant, any result will neither support nor refute the rule that you are testing.
- If you turn over the third card and there is a vowel on the other side, it will support the rule. However, if there is a consonant, it will not tell you anything about the rule (the rule does not say that consonants must be coupled with odd numbers). So, turning over the third card can potentially support your rule, but cannot refute it.
- This is not the case with the fourth card. If the letter on the other side is a consonant, it tells you nothing about the rule. If it is a vowel, though, that would refute the rule.

So the two options that can potentially refute the rule are A and 7. Only in this combination of trials are we actually testing ("falsifying") the hypothesis. It goes in line with the logic of hypothesis testing widely accepted in science. Why do individuals normally give a different pattern of responses? Presumably people are much more attentive to information that can potentially support their expectations and at the same time they tend to ignore information that can potentially contradict their expectations. This preference for potentially supporting evidence is known as confirmation bias.

TOK

What do you know about the falsification principle (Popper, 1959)? What does it mean for a theory to be "falsifiable"? What is the role of falsification in sciences compared to other areas of knowledge?

Congruence bias

Some authors argue that the term "confirmation bias" does not accurately describe these decision-making patterns and offer the term "congruence bias" instead, stressing that the cause of the heuristic is the subject's failure to come up with alternative hypotheses. In other words, the essential thinking behind the congruence heuristic is as follows: "To test a hypothesis, think of a result that would be found if the hypothesis were true and then look for that result (and do not worry about other hypotheses that might yield the same result" (Baron, 2008, p 173).

Congruence bias makes subjects act as if they were trying to obtain positive results (supporting evidence) rather than useful information. This is seen in **Tschirgi (1980)** who gave subjects problems like the following.

John decided to make a cake. When he ran out of some ingredients, this is what he did.

- He used margarine instead of butter for the shortening.
- He used honey instead of sugar for the sweetening.
- He used brown wholewheat flour instead of regular white flour.

The cake turned out great; it was so moist. John thought the reason the cake was so great was the honey. He thought that the type of shortening (butter or margarine) or the type of flour really didn't matter. What should he do to prove his point when he makes the next cake?

Subjects were required to choose from the following options.

A. Use sugar instead of honey

B. Keep the honey, but change everything else

C. Change everythin.

Most subjects chose to keep the honey, but change everything else (option B). Is this the "correct" response? Let's think logically. If John's hypothesis is true, option A will ruin the cake, but support the hypothesis (that is, provide valuable information). Option B in this case will not ruin the cake, but will provide equally valuable information. Option C is a distractor. So it seems like both option A and B potentially are valuable in terms of hypothesis-testing. Interestingly, subjects demonstrated a preference for option B.

In another condition in the study, the description and the options were the same, but subjects were told that: "The cake turned out just terrible. It was so runny". Again, John thought that the reason the cake was so terrible was the honey.

In this condition subjects demonstrated a preference for option A, use sugar instead of honey.

Again, both options A and B are potentially informative in terms of hypothesis-testing, but this time option B ruins the cake, whereas option A does not.

Researchers concluded that subjects in the experiment were pursuing the goal of obtaining a result that would be **"positive" rather than "informative"**.

There seems to be nothing wrong with trying to test a hypothesis and cook a great cake at the same time. We will compare this to another study done by **Wason (1960)** but first do the exercise to try the study for yourself.

Exercise

Ask your friend to read about Wason's experiment below and then construct his or her own experiment by analogy.

Let your friend lead your participation. It will only take five minutes.

Were you able to identify the correct rule?

In this study subjects were given a sequence of numbers (for example, 2, 4, 6) and told that the sequence followed a rule. The task was to discover the rule. To do that, participants could generate additional three-number sequences, and the experimenter told them whether these new sequences followed the rule or not. Usually the subjects would assume that the rule was "numbers ascending by 2", which was reflected in the testing sequences that they produced: 1, 3, 5; 8, 10, 12; 50, 52, 54, and so on. Every time they got positive feedback from the researcher ("yes, this sequence follows the same rule"). The crucial observation here was that subjects rarely questioned their favoured hypothesis and after several trials they usually claimed to have discovered the rule.

ATL skills: Thinking

In what ways does Wason's (1960) rule discovery task resemble Tschirgi's (1980) cake task? Is this essentially the same task?

However, other hypotheses can fit the sequence 2, 4, 6 equally well. In fact, the correct rule in this study was "any ascending numbers". In order to find it out, a typical participant of this study, just like a typical real-world scientist, had to generate alternative hypotheses (for example, 1, 2, 3 or 6, 4, 2), which rarely happened, due to congruence bias. Subjects are looking for positive results that favour their initial expectation rather than informative results. This may be really detrimental when it comes to real science! Unlike making a cake, science requires you to seek out informative results, not positive ones.

TOK

Think about other examples (both from sciences and other areas of knowledge) where congruence bias is a common occurrence.

Illusory correlations and implicit personality theories

The tendency to seek out information that confirms pre-existing beliefs is also seen in illusory correlations and implicit personality theories. An illusory correlation is a belief that two phenomena are connected when in fact they are not. You will come across illusory correlations in the sociocultural approach to behaviour when you study stereotypes, because illusory correlations are often believed to be the mechanism of stereotype formation.

Implicit personality theories are sets of beliefs that you have about the behaviour of others; you predict their behaviour on the basis of those beliefs. For example, you may implicitly believe that all muscular, bald men are dangerous (based on the history of your interactions with them in the past or maybe a number of movies that you have watched), and so you would avoid bald, well-built males in a variety of situations due to your (stereotyped) implicit personality theory.

Discussion

Find out more about the history of using the Rorschach ink-blot test. Present your findings in class.

What do you think about projective tests in general?

What is the role of confirming pre-existing beliefs in the formation of illusory correlations and implicit personality theories? **Chapman and Chapman (1969)** demonstrated this in a sample of practising psychodiagnosticians (N = 32) who used the Rorschach ink-blot test in their practice. They concentrated specifically on diagnosing male homosexuality. Prior research had revealed some Rorschach signs that are statistically associated with male homosexuality and some that are not. Two signs that had been shown to be clinically valid signs of male homosexuality are:

- response on Card IV of "human or animal – contorted, monstrous, or threatening"; examples would be "a horrid beast" or "a giant with shrunken arms"
- response on Card V of an "animalized human or humanized animal"; examples would be "pigeon wearing mittens" or "a woman, dressed as a bat".

However, when asked to recollect their clinical experience and name the Rorschach signs that they had found to be most diagnostic of homosexuality, clinicians failed to mention these two signs and named other signs instead, for example, feminine clothing ("a woman's bra" in Card III), humans with confused or uncertain sex, male or female genitalia.

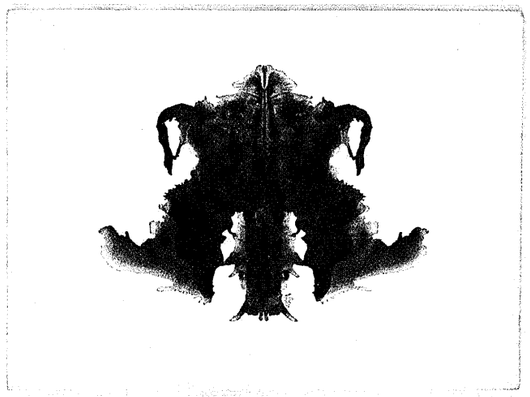

▲ Figure 3.20 Rorschach Card IV

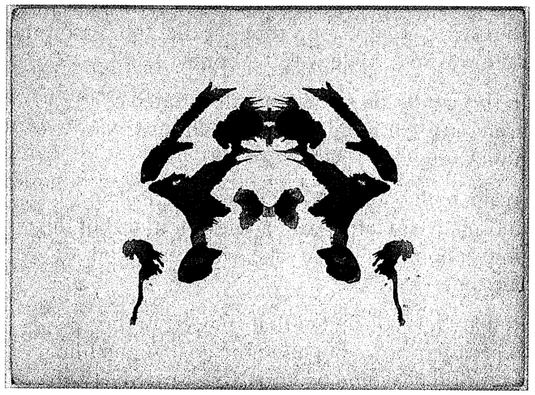

▲ Figure 3.21 Rorschach Card V

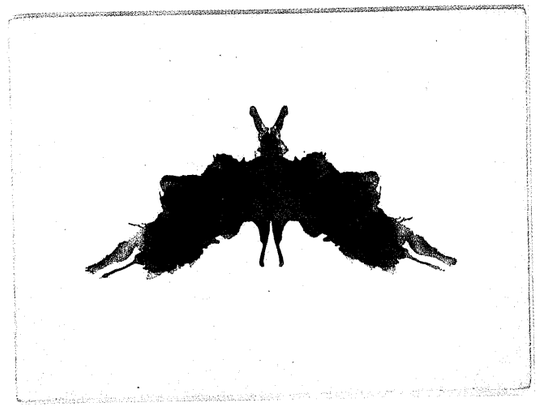

▲ Figure 3.22 Rorschach Card III

All these signs had a strong verbal associative connection to homosexuality. When asked to rate the associative similarity of a sign with homosexuality, clinicians rated the similarity as high for the popular (and invalid) signs and low for the valid (and unpopular) ones. For example, contrary to statistical evidence, they said that seeing "a woman's bra" in Card III was a sign of homosexuality, and at the same time they failed to recognize seeing "a horrid beast" in Card V as one.

Chapman and Chapman (1969) also studied whether naive observers would make the same errors as the clinicians did. Participants were students in an introductory psychology course. The fabricated clinical materials consisted of 30 Rorschach cards, on each of which there was one response about the ink blot and two statements about the patient who (allegedly) gave this response. For example, the typical card would show an ink blot and three statements:

- Response: "A pigeon wearing mittens"
- Statement: A man who said he has sexual feelings towards other men
- Statement: A man who said he feels sad and depressed much of the time.

The response was taken either from a valid diagnostic category (for example, "A giant with shrunken arms" for Card IV) or an invalid one (for example, "a woman's bra" for Card III). The two statements were taken from a pool of four symptoms—homosexuality (in 1969 when the study was conducted homosexuality was still considered to be a mental disorder), depression, paranoia, inferiority complex. The combinations of responses and statements were manipulated to be completely random, that is, there was no statistical relation between homosexuality and the frequency of any of the responses. However, when asked to formulate a rule after being exposed to all 30 cards (what kind of responses are most common for homosexuals?) participants readily named the invalid ones (for example, seeing "a woman's bra" in Card III). So, naive participants arrived at the same results, and used the same justifications, as experienced clinicians, and yet in neither of the groups were the results valid! Participants seemed to have a set of prior beliefs (probably based on common sense) and they were selectively interpreting available data to support, but not contradict, those beliefs.

Even more strikingly, when in follow-up experiments Chapman and Chapman manipulated the valid signs to **actually correlate** with homosexuality (in most of the cases mentioning homosexuality on the card was coupled with a valid sign such as seeing a "horrid beast" in Card IV), this had practically **no effect** on the subjects' conclusions. They still failed to see the connection between homosexuality and the valid signs, and continued to see a connection with the invalid ones (such as seeing a "woman's bra" in Card III). Illusory correlations based on prior beliefs turn out to be quite stable and resistant to change even in the presence of counterevidence.

ATL skills: Communication

Do you think illusory correlations and implicit personality theories can affect psychological and medical practices on a wider scale? Would you go as far as saying that they are inevitable?

The tendency to avoid the mental stress of holding inconsistent cognitions

Another source of cognitive biases and system I thinking is the tendency to avoid the mental stress of holding inconsistent cognitions.

This was extensively demonstrated by **Leon Festinger** in his **theory of cognitive dissonance** as well as in supporting research. Cognitive

dissonance is the mental stress caused by the inconsistency between:

- two (or more) contradictory beliefs or ideas
- one's action and one's belief
- new information and existing beliefs.

According to the theory, an individual who experiences such inconsistency ("dissonance") feels stressed and uncomfortable, and is driven by a desire to **reduce dissonance**. In the action versus belief example, there are only two logical ways to reduce dissonance: change your behaviour or change your belief. It would be logical to assume, from the common sense perspective, that "beliefs drive behaviour". However, many research studies have actually shown that "behaviour drives beliefs" to a much larger extent.

One way to empirically test the predictions of the theory of cognitive dissonance is through observation of groups of people who strongly believe in something but new evidence runs contrary to their beliefs. This research framework is known as the "**belief disconfirmation paradigm**". A well-known example, and the birth of the theory of cognitive dissonance, is Leon Festinger's book *When Prophecy Fails* (1956). In this he describes a small UFO cult in Chicago called the Seekers, who believed in an imminent apocalypse and took strong actions to support their belief. The leader of the cult, a self-proclaimed prophet, "received messages" that the world would end at midnight on 21 December 1954. They also believed that they were selected to be the survivors and start a new civilization, and for this, a spacecraft was coming to pick them up and they had to prepare for departure. As part of the preparation they had given up their jobs, spouses and possessions, which shows how committed they were to the belief. Festinger infiltrated the group to conduct participant observations and collected valuable information on what happens if such strong beliefs are disconfirmed by evidence.

After the prophecy proved to be false (the space shuttle did not pick them up at midnight) the group spent several hours in a state of tension which is very well conveyed in Festinger's detailed observation report. At 4.45am the leader went to a separate room and received another "message" which said that the apocalypse had been called off because the members of the religious group saved the Earth by their pure thoughts and efforts: "The little group, sitting all night long, had spread so much light that God had saved the world from destruction" (Festinger, Riecken and Schachter, 1956, p 169, *When Prophecy Fails: A Social and Psychological Study of a Modern Group that Predicted the Destruction of the World*). Interestingly, although the group had been notoriously closed to the world outside and had never given interviews, on the morning of the following day they began a campaign to spread the message as far as possible, calling newspapers and setting up interviews.

This shows how, when new information disconfirms existing beliefs, **the belief itself may be twisted**, and social support may be sought to further support this new belief; if many people start sharing the belief, then it becomes more justified. This also shows how people change their beliefs when it is impossible to change (take back) their behaviour.

Another research framework that is often used to test the theory of cognitive dissonance is known as the **induced-compliance paradigm**. An example is the study of **Freedman and Fraser (1966)** in which participants from the experimental group were asked to sign a petition on the issue of safe driving (participants from the control group were not asked to do anything). Two weeks later, all participants were asked to put a large sign on their front lawn saying "Drive Carefully" (a larger, more substantial request). Results showed that fewer than 20% of the control group agreed to put the large sign on their lawn, while over 55% of subjects in the experimental group agreed to this larger request.

This shows how once a behaviour is demonstrated ("I signed a petition for safe driving"), a behaviour–belief discrepancy may occur ("If I don't want to put this large sign on my lawn, why did I sign the petition two weeks ago?"). To avoid psychological discomfort caused by this dissonance, subjects would try to adjust their initial beliefs ("Maybe I feel strongly about safe driving, after all, and I support their cause").

Many effective persuasion techniques are based on the principles of cognitive dissonance.

Exercise

Explore some other well-known heuristics and biases, as well as a range of their potential practical applications:

Dan Ariely's book *Predictably Irrational*: https://tinyurl.com/k683ou4

Daniel Kahneman's book *Thinking, Fast and Slow*: https://tinyurl.com/hx6w569

Laurie Santos's TED Talk "A monkey economy as irrational as ours" (in which she compares economic decisions in humans and monkeys and finds some striking similarities): https://www.ted.com/talks/laurie_santos

Make a list of the most commonly mentioned cognitive biases and make it a rule for the next week to observe and take a note of real-life examples of heuristics and cognitive biases in people's decisions.