

(e) If this simple model fails, construct further models which might involve one or more of the following possibilities:

(i) Relax assumptions of full information, probably following recognized procedure in game theory.

(ii) Relax the assumption of self-interest, introducing richer utility/preference functions (e.g., altruism, malice, indifference, envy, etc.).

(iii) Relax the assumption of objectively calculated resources introducing subjective beliefs about resources.

(iv) Relax the assumption of a set of feasible actions, calculated objectively, introducing subjective beliefs about what is feasible.

(v) Re-compute the structure of interdependencies.

When either (i), (ii), (iii), or (iv) are seen to hold, this will inevitably prompt a further question as to why it should be so. Answers to these questions might also be couched in terms of some of the RCT precepts outlined above (e.g., strategic limitations of information). It is only if this very general framework fails to provide the answers we seek that we should then reach for an alternative theoretical framework.

See also: Action, Theories of Social; Altruism and Self-interest; Bounded Rationality; Coleman, James Samuel (1926–95); Game Theory; Interests, Sociological Analysis of; Macrosociology–Microsociology; Methodological Individualism in Sociology; Rational Choice Explanation: Philosophical Aspects; Rational Choice in Politics; Rational Choice Theory: Cultural Concerns; Sociology: Overview; Theory: Sociological; Traditions in Sociology

Bibliography

- Abell P (ed.) 1991 *Rational Choice Theory*. Elgar, Aldershot, UK
- Abell P 1996 Sociological theory and rational choice theory. In: Turner B S (ed.) *Social Theory*. Blackwell, Oxford, UK
- Boudon R 1989 Subjective rationality and the exploration of social behaviour. *Rationality and Society* 12: 173–96
- Coleman J S, Fararo T J 1992 *Rational Choice Theory: Advocacy and Critique*. Sage, Newbury Park, CA
- Moser P K (ed.) 1990 *Rationality in Action: Contemporary Approaches*. Cambridge University Press, Cambridge, UK

P. Abell

Rational Theory of Cognition in Psychology

1. People Appear Irrational in Cognitive Experiments

Michael Watkins has said that a cognitive theory ‘is a bit like someone else’s toothbrush—it is fine for that

individual’s use, but for the rest of us ... well, we would just rather not, thank you’ (Watkins 1984, p. 86). To say whether people are behaving rationally requires a definition of what it means to behave rationally, and like a toothbrush everyone has their own. For the purposes of this article, rational behavior will be defined as follows: to behave rationally in some context is to display behavior that corresponds to a normative standard of behavior for that context. The choice of the normative standard determines what constitutes rational behavior. For the economist, behaving rationally involves maximizing utility; for the logician it is following the deductive rules of logic; and for the (Bayesian) statistician it is acting according to Bayes’s rule. One need not look far to find that people do not reason rationally with respect to a variety of normative standards. There are entire literatures on how and why people violate the rules of deductive logic and statistical inference. In the Wason (1968) card task, central to one such literature, people are given a rule of the form ‘if p, then q,’ such as ‘if there is a vowel printed on one side of a card, then an even number is printed on the other.’ The participants are next presented with a set of cards, such as A, K, 2, 7. Their task is to choose only those cards that need to be flipped to check whether the rule holds. In this example, only the A-card (p) and the 7-card (\sim q) need to be checked. An odd number on the opposite side of the A-card would clearly violate the rule as would a vowel on the other side of the 7-card (\sim q). The rule says nothing about what is on the opposite side of a consonant, so flipping the K-card (\sim p) does not test the rule. Also, flipping the 2-card (q) cannot disconfirm the rule, because the rule does not restrict what is on the opposite side of an even-numbered card. In general people perform terribly on this task, or at least their performance does not appear to be rational with respect to the rules of deduction. Oaksford and Chater (1998) surveyed 13 studies, covering a variety of Wason tasks. The proportions of people who flip the p(A), \sim p(K), q(2), \sim q(7) cards were 89, 16, 62, and 25 percent respectively. That is, people rarely flip the \sim q card, which they should flip, and frequently flip the q card, which they need not do.

People are similarly deficient in drawing inferences from statistical information. Consider the following example that requires people to interpret the results of a positive colon cancer test (Gigerenzer 1998). Suppose that 0.3 percent of people have colon cancer, there is a 50 percent chance that a colon cancer test will detect a cancer (a hit), and a 3 percent chance that it will indicate that there is cancer when there is none (a false positive). What is the probability that a person with a positive test has cancer? When Gigerenzer posed this problem to physicians their median estimate was 47 percent. The normative answer is more like 4.8 percent (see Table 1). This answer is obtained by applying Bayes’s rule to the statistics given. Bayes’s rule shows the correct way to make statistical inferences. Bayes’s

Table 1

Bayes's rule applied to Gigerenzer's (1998) cancer example

This shows the probability of some hypothesis, H_j being true given some data, D . Bayes's rule takes the following form:

$$p(H_j|D) = \frac{p(H_j) p(D|H_j)}{\sum_i p(H_i) p(D|H_i)}$$

where $p(H_j)$ is the belief that the hypothesis H_j is true prior to observing the data, D , and $p(D|H_j)$ is the probability of observing the data, given the hypothesis is true. The denominator is the probability of observing the data. It assumes that there are i different hypotheses, one of which is true. We can apply Bayes's rule to Gigerenzer's (1998) cancer example. The two hypotheses are that the patient has cancer, or that he or she does not (\sim cancer). The base rate, or prior, probability of having colon cancer is 0.3 percent. The hit rate of the test is 50 percent, and the false alarm rate is 3 percent. Applying Bayes's rule to this example yields an estimate of having cancer to be 4.8 percent, a much more reassuring estimate than given by the doctors.

$$p(\text{cancer} | + \text{test}) = \frac{p(\text{cancer}) * p(+ \text{test} | \text{cancer})}{p(\text{cancer}) * p(+ \text{test} | \text{cancer}) + p(\sim \text{cancer}) * p(+ \text{test} | \sim \text{cancer})} = \frac{.003 * .5}{.003 * .5 + .997 * .03} = .0477$$

rule implies that given the rarity of cancer, a positive test is more likely to follow a false positive than the actual detection of cancer. As with the Wason (1968) task, subjects' performance on statistical tasks like this clearly shows that they are behaving irrationally. The standard explanation for this is that people tend to ignore base rates; in this example they fail to take into account that cancer is uncommon. That is, only 0.3 percent of people have cancer.

2. People Employ Ecologically Valid Strategies in the Laboratory

Performance on Wason (1968) card tasks, Gigerenzer's (1998) statistical tasks, and other experiments demonstrate that people are, in fact, irrational, when measured against accepted standard inferential rules. There has been a number of proposals to describe what people are doing (e.g., Tversky and Kahneman 1974). The idea behind these approaches is that people are applying heuristics, rules of thumb that tend to be successful in their daily lives, to these laboratory situations where they do not apply. Anderson (1990) has taken this idea one step further, arguing that people are not just bringing real-world heuristics into the laboratory, but rather are employing optimal solutions to certain problems faced in their natural environments. However, behavior that is rational in natural environments may not necessarily be rational in the peculiar environments that experimental psychologists concoct. Some of the roots of this idea are discussed next.

3. David Marr's Levels of Explanation

Anderson's motivations closely follow the arguments that David Marr laid out in his influential book *Vision*.

Marr (1982) argues that fully understanding an information processing system requires considering the system from multiple levels. Marr was interested in the visual system, which he argues evolved to 'tell(s) us about shape and space and spatial arrangement' (Marr 1982, p. 36). The distinction that Marr makes among the levels of explanation can be more readily understood by referring to a far simpler information processing system than vision—a clock. What does a clock do? It indexes time. How does it do it? It achieves this goal by measuring the passage of time by incrementing counters at fixed intervals. Marr says that descriptions like this are at the computational level.

The next level down, called the representation and algorithm level, describes how the goals should be achieved. In the case of a clock it would map the cycles of an oscillator into seconds, minutes, and hours. Descriptions at this level of analysis require a specification of the representations and the algorithms that operate on them to achieve the goals specified at the computational level. Many combinations of representation and algorithms can achieve these goals; a 12 or 24 hour clock can index time. The choice of representation does constrain the choice of algorithm, such as how seconds are rolled over into minutes, and minutes into hours (e.g., what happens at 12:59). Further, not all representation and algorithms are equivalent; some computations may be simpler with one representation than another. Calculating the duration of a trip is simpler when the train leaves at 10:00 and arrives at 14:00 than when it leaves at 10:00 a.m. and arrives at 2:00 p.m.

The lowest level of description in Marr's hierarchy, the hardware implementation level, is concerned with describing the physical entities that carry out the computations of the representation and algorithm level. Here, the current time could be represented by the configuration of hands on a clock face, the coordination of seconds and minutes handled by brass

wheels, the oscillator could be realized as a pendulum. Marr's point is that if you stumbled onto a morass of gears, you would be better able to make sense of these gears if you knew that they were part of a clock as opposed to a cash register, or a sewing machine. Similarly, one is going to be better able to understand people as information processing systems, if one understands what those systems are trying to accomplish. As Marr puts it, 'trying to understand perception by studying only neurons is like trying to understand bird flight by studying only feathers: it just cannot be done. In order to understand bird flight, we have to understand aerodynamics; only then do the structures of feathers and the different shapes of birds' wings make sense' (Marr 1982, p. 27).

4. John Anderson's Principle of Rationality

Marr (1982) demonstrated the utility of approaching the problem of vision from multiple levels, and particularly from the computational level. Based on Marr's success with vision, Anderson (1990) hypothesized that the approach might work well for higher-level cognitive processes. Anderson argued that most theorizing in cognitive psychology concerns representations and the processes that act on them, corresponding to Marr's representation and algorithm level. Cognitive neuroscience and neuroscience focus on the hardware implementation level. Relatively little theorizing, he points out, occurs at Marr's computational level. Anderson (1990) explored whether it would be useful to think about higher level cognition, such as categorization (Anderson and Matessa 1992) and memory (Anderson and Milson 1989, Anderson and Schooler 2000) at the rational level, analogous to Marr's computational level. His principle of rationality is that 'the cognitive system operates at all times to optimize the adaptation of the behavior of the organism' (Anderson 1990, p. 28). Anderson renamed Marr's computational level the rational level, because the computational level sounds like it should be describing the algorithm and representation level. His choice of the term rational also serves as an allusion to economic theorizing that often takes place at Marr's computational level. Economists focus more on the decisions agents make, rather than on the processes involved in coming to those decisions. In the economic realm, it is easy to appreciate that it is rational for firms to maximize profits (or, as in the case of Ben and Jerry's, the famous American premium ice-cream maker, to maximize a combination of profits and social good). For the cognitive system, the currency is less clear. Thus, the critical step in what Anderson calls a 'rational analysis' is to figure out what quantity the cognitive system might be optimizing, and to make predictions based on this about how people will behave in particular experimental tasks.

5. Oaksford and Chater's Rational Analysis of the Wason Card Task

There would seem to be a weak case for the rationality of the cognitive system in light of people's irrational behavior on the Wason (1968) card task and Gigerenzer's (1998) cancer problem. However, Oaksford and Chater's (1996) rational analysis of the Wason card task shows that typical performance on the Wason, while violating traditional notions of rationality, is indeed quite rational when seen in a broader context of how people seek information in the world.

Consider the following situations that parents may face. In one situation, hearing an utterance from the baby he was caring for, a young father might ask 'was that just babbling?' or if the baby says 'baba,' then the baby wants a bottle. If the mother found her husband giving the baby a bottle, and if she wanted to know whether 'baba' was serving as the baby's signal it wanted a bottle, would it be rational for the mother to ask whether the baby had said 'baba' before the feeding? Now imagine 16 years have passed, and the mother is teaching her son to drive. She wonders whether he knows that when the oil light is on, he should stop the car. Would it be rational for the mother to ask whether the oil light is on? It would seem quite natural for the young mother to ask about what the baby had said, and for the middle-aged mother to remain silent. Asking in the first case would be like flipping the 2-card, irrational with respect to deductive inference. Asking cannot disconfirm the hypothesis that the baby knows how to use the word, since logically many circumstances (e.g., fussing) can suggest an infant is hungry. In contrast, asking about the oil light would be like flipping the 7-card, rational with respect to the rules of logical inference. If the mother found that the oil light was on, this would be a violation of the rule. Our intuitions are at odds with logic.

Perhaps our intuitions are correct that the young mother is rational to ask about what the baby said, and the older mother is rational in remaining silent. The essential difference between the two situations is the amount of information the answer to the question is likely to yield. Though asking whether the baby said 'baba' could not provide a definitive answer, knowing what the baby said tells something about whether the baby knows the word. In contrast, asking about the oil light, though potentially highly informative in the unlikely event that the oil light is on, will tell nothing in the more likely event that the oil light is off.

Oaksford and Chater's (1996) analysis of the Wason task formalizes these intuitions. They assume that people apply experimental, information-seeking strategies to deductive tasks like the Wason. Though the details of their mathematical analysis are beyond the scope of this article, the flavor of it can be given here. As is true for nearly all rational analyses, they assume

that people are behaving as if they are following Bayes's rule. They suggest that people are not treating the task as a test of logic, but rather are attempting to gauge the causal relation between two events. More specifically, they assume what people are really trying to do in deductive tasks is decide between two hypotheses: when p occurs (e.g., baby says 'baba'), q must follow (e.g., baby wants a bottle), or the alternative hypothesis that event p is independent of event q . Sometimes the baby says 'baba,' and sometimes the baby is hungry, and it is only by chance that the baby says 'baba' when it is hungry. In the case of the Wason task described earlier, the competing hypotheses are that an even number depends on a vowel or the alternative that evens and vowels are independent of each other. The question, then, is which card will provide the most evidence in terms of discriminating between these two rival hypotheses. Since people do not know in advance how an experiment (i.e., flipping a card) is going to turn out, they make their decisions based on how much information they expect to gain.

A critical difference between the parenting examples and the Wason (1968) task is that people have experience with how children learn words, but relatively little experience with numbered and lettered cards. Lacking any relevant experience about the cards, Oaksford and Chater (1996) assume that people treat them as if they are typical of causal relations they have seen in the past. When their model of information-seeking is combined with the assumption that causally related events are relatively rare, it predicts the observed subject preferences for flipping cards, namely p (e.g., A-card) is chosen more than q (e.g., 2-card), q more than $\sim p$ (e.g., K-card), and $\sim p$ more than $\sim q$ (e.g., 7-card). So in the Wason task we have a situation where the behavior is irrational with respect to the laws of deduction, but can be understood to be rational in the context of how people seek information.

6. *People are Rational When Ecologically Valid Strategies are Appropriate*

Gigerenzer's (1998) colon cancer experiment demonstrated that people are irrational with respect to proper Bayesian inference. This result seems to be at odds with Anderson's rational hypothesis, as Bayesian analyses underlie nearly all rational analyses. Since the problem was laid out perfectly for Bayesian inference, we would expect people to perform well. Gigerenzer argues that people perform poorly, because the problem format is all wrong. People did not evolve, he argues, to process probabilities, but rather to draw inferences based on their direct experience. His prediction is that people should do well on statistical tasks when the data are presented more naturally in terms of raw frequencies. Consider again the colon cancer problem, but this time in raw frequencies. Thirty out

of every 10,000 people have colon cancer. Of these 30, 15 will test positive. Of the remaining 9,970 people without cancer, 300 will still test positive. What proportion of people who test positive will actually have cancer? Now it is clear that only 15 of the 315 (or 4.45 percent) of those who test positive will have cancer. When the problem was presented this way to another group of doctors, 67 percent answered correctly, as compared to 4 percent when the data were presented in terms of probabilities. When the problem presentation is consistent with how we experience events in the world, people's behavior corresponds to the prescriptions of a normative standard, a Bayesian inference. As with the Wason task, people appear to be behaving irrationally, when the experimental task and conditions conflict with the natural environment.

7. *In Practice, Rational Analyses are Bounded*

The discussions of the Wason (1968) task and Gigerenzer's (1998) task were couched strictly at the rational (or computational) level. In practice, a rational analysis cannot focus strictly on the rational level, but must also consider the algorithm and representation level. The reason for this is that sometimes the rational solution requires calculations that would be physically impossible for any system to perform. In such circumstances a mapping needs to be made from the results of the rational level into algorithms and representations that approximate the computations called for by the rational analysis. It is at the algorithm and representation level that compromises have to be made. In particular, assumptions have to be made about processing limitations. This ties in with Simon's (1956) notion of bounded rationality. That is, people are rational within the constraints of their ability to process information. For example, Simon (1990) pointed out that ignoring processing limitations suggests that knowledge of the rules of chess should lead to perfect play. If a person or machine had an infinite amount of time to contemplate a chess move, and the potential countermoves of the opponent to that move, and the countermoves to all the opponent's potential countermoves, *ad infinitum*, then the person could select an initial move that would inevitably lead them to check mate. The problem is that more board positions would have to be considered than there are molecules in the universe (Simon 1990). Instead, people and chess programs use heuristics to mimic the behavior of this rational strategy. Similarly, because of processing limitations, most rational analyses are approximations to the optimal solution.

There are some games, however, for which it is possible to know the game completely. For example, many adults know the moves and countermoves in tic-tac-toe to guarantee that they will at least not lose. Like tic-tac-toe, the Wason (1968) task is atypical. Thus the number of potential hypotheses and experi-

ments raised by Oaksford and Chater's (1996) analysis of the task is relatively small compared to the number raised by problems people often face in their daily and work lives. For example, consider again the problem of interpreting what the baby means when it says 'baba.' The potential number of hypotheses for what 'baba' means certainly exceeds the number of configurations of a chessboard. If Oaksford and Chater's information-seeking strategy were applied to the problem of vocabulary acquisition, where large numbers of hypotheses and experiments are potentially relevant, they would need to employ heuristics and build processing bounds into their system as well.

8. *Whether People are Rational Depends on Your Perspective*

This article demonstrated that there is no definitive answer to the question of whether people behave rationally. For example, the Wason (1968) task clearly demonstrates that people behave irrationally with respect to logic, while Oaksford and Chater's (1996) analysis shows that the behavior is rational when the context is broadened to encompass a person's "normal" life conditions' (Brunswick 1943, p. 259). In short, people can be seen to behave rationally with respect to the environment, but appear to be operating irrationally with respect to a particular task, especially when the task that subjects are performing differs from the one intended by the experimenter.

See also: Functionalism, History of

Bibliography

- Anderson J R, Matessa M 1992 Explorations of an incremental, Bayesian algorithm for categorization. *Machine Learning* **9**: 275–308
- Anderson J R, Milson R 1989 Human memory: An adaptive perspective. *Psychological Review* **96**: 703–19
- Anderson J R, Schooler L J 2000 The adaptive nature of memory. In: Tulving E, Craik F I M (eds.) *Handbook of Memory*. Oxford University Press, New York
- Brunswick E 1943 Organismic achievement and environmental probability. *Psychological Review* **50**: 255–72
- Gigerenzer G 1998 Ecological intelligence: An adaptation for frequencies. In: Cummins D D, Allen C (eds.) *The Evolution of Mind*. Academic Press, San Diego, CA
- Gigerenzer G, Todd P M, the ABC Research Group 1999 *Simple Heuristics that Make Us Smart*. Oxford University Press, New York
- Manktelow K 1999 *Reasoning and Thinking*. Psychology Press, Hove, UK
- Marr D 1982 *Vision*. Freeman, San Francisco
- Oaksford M, Chater N 1996 Rational explanation of the selection task. *Psychological Review* **103**: 381–91
- Oaksford M, Chater N 1998a *Rationality in an Uncertain World*. Psychology Press, Hove, UK
- Oaksford M, Chater N (eds.) 1998b *Rational Models of Cognition*. Oxford University Press, New York
- Simon H A 1956 Rational choice and the structure of the environment. *Psychological Review* **63**: 129–38

- Simon H A 1990 Invariants of human behavior. In: Rozenweig M R, Porter L W (eds.) *Annual Review of Psychology*. Annual Reviews, Palo Alto, CA
- Tversky A, Kahneman D 1974 Judgement under uncertainty: Heuristics and biases. *Science* **185**: 1124–31
- Wason P C 1968 Reasoning about a rule. *Quarterly Journal of Experimental Psychology* **20**: 273–81
- Watkins M J 1984 Models as toothbrushes. *The Behavioral and Brain Science* **7**: 86

L. Schooler

Rationalism

When philosophy consisted of set piece battles between grand schools of thought (such as realism, scepticism, or monism), 'rationalism' referred to the belief that human beings disposed of a faculty called 'reason' which gave them access to the structure of reality. Man, wrote Aristotle (384–322 BC), is a rational animal, but Aristotle recognized that reason had to be trained. In most human beings, the understanding was distorted by passion. Aristotle's teacher Plato (427–347 BC) had argued in *The Republic* and other dialogues that the world we experience was a confused copy of a world of forms or ideas which could be discovered by philosophical inquiry. His parable of the cave (*Republic*, Book VII) is the founding image of rationalism. In its classical Greek version, rationalism assumed that we might understand the structure of the universe by the power of reason, an assumption connected with the fact that rationalism began its philosophical career as a generalization of the procedures of mathematics. In his dialogue the *Meno*. Plato had argued that knowledge is accessible to rational inquiry independently of experience. Later, the Stoics argued that moral knowledge of the laws of nature is available to any rational creature who looks into himself.

Medieval philosophy revived Greek ideas and cross-fertilized them with Christian doctrine. St. Thomas Aquinas (1225–1274 AD) combined reason and revelation by taking his account of nature from Aristotle, and his account of higher things (or 'supernature') from Christian revelation. This synthesis, however, soon began to fall apart, and in modern times rationalism has commonly been taken to refer to rejection of Christian revelation because of rational criticism.

1. *Modern Rationalism*

This was not, however, the view taken by the great rationalist philosophers of the seventeenth century, for whom science was the model of understanding. Descartes invoked God's veracity as a foundation of his system, while Spinoza was a monist who argued

12775