bounded rationality There are three interpretations of the concept of bounded rationality (see RATIONALITY). The first defines it as a deviation of unbounded rationality due to external constraints, such as information and deliberation costs. This view, called *optimization under constraints*, is a standard modelling technique in economics and beyond (Arrow 2004). The second interpretation defines bounded rationality as a deviation from unbounded rationality because of internal constraints, notably limited information-processing capacities and lack of willpower. This *cognitive illusions* or *heuristics and biases* approach, as it is variously called, attributes the deviations to cognitive limitations (Kahneman 2003). The two interpretations appear diametrically opposed. By looking at the constraints in the world, the first emphasizes the rationality of human action; by focusing on the constraints in the mind, the second perceives irrationality.

The third interpretation is known as *ecological rationality*; it focuses on how *heuristics succeed in the real world rather than in comparison with unbounded rationality (Gigerenzer et al. 1999). The difference between the three interpretations can be explained in terms of two ideals underlying the concept of unbounded rationality: *optimization* and *universality*. Optimization under constraints retains both: It models how to find the optimal (best) action given constraints based on a universal calculus (such as maximizing expected utility). The cognitive illusions programme also retains the ideal of optimality, but attributes deviations to cognitive limitations rather than to optimization under constraints. Finding the optimal action is, however, usually impossible when it comes to real-world problems. For instance, the optimal sequence of moves in chess cannot be found, neither by Deep Blue nor Kasparov. Most interesting problems are computationally intractable (or NP-hard) or do not allow optimization for other reasons, such as when multiple goals compete, when the criterion cannot be precisely measured (with concepts such as happiness), and when the problem is ill-defined (finding the best wife or husband). Optimization under constraints deals with this problem by transforming complex real-world problems into 'small worlds' and introducing convenient mathematical assumptions that allow optimization. The hope is that the optimal behaviour determined for the small world might also be optimal in the messy real world.

In contrast, the programme of ecological rationality asks how real people make decisions when optimization is out of reach. It replaces optimizing by satisficing (finding an action that is good enough), and universality by modularity. The programme is threefold: (1) to describe the heuristics in the 'adaptive toolbox', their building blocks, and the evolved capacities they exploit; (2) to determine the social and physical environments in which a given heuristic is better than other strategies, that is, their ecological rationality; and (3) to design heuristics and environments in applied contexts, such as medicine, law, business, and engineering. Examples of heuristics include the recognition heuristic (e.g. relying on brand name recognition in consumer choice), one-reason decision-making (basing a complex decision on one good reason only), equality (dividing your assets equally, as in investment and food sharing), peer group imitation, advice taking (learning from the experience of others), and default (e.g. abiding by a legal default for organ donation, whether opt-in or opt-out). Examples of environmental structures that support one-reason decision-making include moderate to high uncertainty, redundancy, and variability of cue weights. In these situations, relying solely on one good reason and ignoring the rest can result in more accurate predictions than complex statistical methods, including multiple regression, neural networks, and classification and regression trees—while being faster and entailing less information cost. In an uncertain world, an intelligent mind needs to ignore part of the information available rather than combining everything. This insight poses a challenge to the ideal that more information and computation is always better; a beneficial amount of cognitive limitation can in fact be adaptive in the sense that we would fare worse without it (Hertwig and Todd 2003).

What is the role of affect in bounded rationality? A supernatural being with perfect knowledge, approximated by the ideal of unbounded rationality, can dispense with emotions. If you can predict the behaviour of others perfectly, then hope, trust, joy, surprise, or fear are redundant. Similarly, these feelings play little role in optimization under constraints. In the cognitive illusions programme, affect has been treated in the same way as cognitive limitations. The study of ecological rationality, in contrast, investigates the situations in which a given emotion or intuition would be good enough for making important decisions (see DECISION-MAKING) (Gigerenzer 2007). For instance, some emotions can function similarly to but more efficiently than the cognitive building blocks of heuristics. Love can stop the search for partners more effectively than an aspiration level, and disgust can limit the choice set more consistently than a consideration set based on weighting of features. The challenge ahead is to incorporate emotions as search or stopping rules into the cognitive models of heuristics, and to analyse the environmental structures that make emotions rational.

GERD GIGERENZER

Gigerenzer, G. and Selten, R. (eds) (2001). *Bounded rationality: the adaptive toolbox*. Cambridge, MA: MIT Press.
Payne, J.W., Bettman, J.R., and Johnson, E.J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.