Gobet, F., & Lane, P. C. R. (2012). Bounded rationality and learning. In N. M. Seel (Ed.), Encyclopedia of the sciences of learning. New York, NY: Springer.

BOUNDED RATIONALITY AND LEARNING

Fernand Gobet

Department of Psychology, Brunel University

Uxbridge

United Kingdom

fernand.gobet@brunel.ac.uk

Peter C. R. Lane

School of Computer Science, University of Hertfordshire

Hatfield

United Kingdom

p.c.lane@herts.ac.uk

Synonyms

Procedural rationality

Definition

Bounded rationality is a term proposed by Nobel Prize winner Herbert A. Simon (1916 – 2001) to emphasize that a decision maker's rational choice is affected by cognitive limitations, in particular limitations in knowledge and limitations in computational capacity. Limitations in knowledge relate to what a decision maker knows about their domain; how long have they been learning, and from which sources? Limitations in computational capacity refer to the ability to identify all relevant factors in a problem, work out their consequences, and come to a conclusion in a reasonable period of time. Psychological explanations for the latter include: small span of attention, limited capacity of short-term memory, slow learning rates, and slow rate of searching problem spaces. All these limitations are particularly problematic given the limited time available for most decisions. Learning is an important way in which the impact of these limitations can be alleviated, at least to some extent.

Theoretical Background

The seeds of Simon's conception can be found in *Administrative Behavior* (Simon, 1947), a book derived from his PhD thesis, in which he showed that the type of unbounded or global rationality assumed by classic economics and statistical decision theory could not account for behavior in organizations. Bounded

rationality, and its impact on human decision-making, guided Simon's research in many domains, including cognitive psychology, artificial intelligence, and the use of information technology in organizations.

Simon introduced the concept of bounded rationality in his analysis of organizational behavior (Simon, 1947), although he was clear that his conclusions applied to the social sciences in general, including in particular economics and psychology. His interest centered on the following question: Can the kind of rationality assumed in classic economics—unlimited access to the relevant information and unlimited resources to carry out the computations necessary for maximizing utility—reflect the way executives make decisions in business, government, and other organizations? Based on a critical analysis of the extant literature, Simon concluded that this was not the case, for two main reasons. First, access to information is strictly limited, and indeed involves a cost. Thus, lacking full information, executives and other decision-makers cannot list all possible alternatives, compute their utility, and choose the optimal solution, as required by unbounded rationality. Second, the human cognitive system is characterized by a number of rather strong limitations. There is a limit of attention: the amount of information which can be perceived, read or simply absorbed at any one time is restricted by constraints, such as the size of our visual field, the rate at which we can read, and the amount of information which can be communicated verbally. Further limits include the capacity of short-term memory, which limits the number of ideas we can hold in mind at any one time, and also the rate with which information can be stored in long-term memory. Given these limits, attention can normally focus on only one thing at a time, and the search for solutions to problems is carried out essentially serially. Thus, even assuming unlimited access to the information required to make a decision, these cognitive limitations make it impossible for humans to carry out the computations necessary for making optimal decisions, except in very simple cases. Rather, executives and humans in general display bounded rationality, that is, they make decisions that are rational given the limited information and computational resources available. They do not try to find an elusive optimal solution, but they *satisfice* and stop their search once a good-enough decision has been met. In this context, satisficing means that a solution to a problem is accepted once all criteria for comparing alternatives are above a given threshold. These criteria are set *a priori* but can then be revised as search progresses and aspiration levels are modified, either downwards or upwards.

In later writings (e.g., Simon, 1997), Simon also used the terms *substantial rationality* and *procedural rationality*, which closely correspond to global rationality and bounded rationality, respectively. Substantial rationality refers to the type of rationality used in economics, where the emphasis is not only on maximizing utility but also on analyzing the situation rather than the decision maker. With procedural rationality, the emphasis is reversed: the interest is about the processes leading to a decision, and the focus is on the decision maker rather than on the situation. Thus, procedural rationality matches the type of rationality studied in psychology.

Given this rather grim description of human cognitive abilities, where the emphasis seems to lie primarily on shortcomings, one is legitimately entitled to wonder how our species was able to produce such developments as the invention of calculus and the discovery of DNA. Simon's answer is that humans (partly) sidestep the limits imposed by bounded rationality by storing previous solutions in long-term memory (through learning), by using powerful heuristics (rules of thumb), and by combining their forces in efficient and typically hierarchical organizations, where decision-making control can be centralized or decentralized as a function of the demands of the internal and external environments. They can also use the environment as an external memory or even processor; for example, by order of sophistication: a cue, such as a road sign, to remind one to carry out an action, paper on which to write notes, and computer memory.

Two major types of decisions can be identified: *programmed decisions* and *non-programmed decisions*. Programmed decisions are well-specified sequences of activities, those repetitive decisions that are routine and well-learned and for which clear-cut procedures have been developed. Thus, it is not necessary to devise new solutions where such decisions must be made. Programmed decisions can be made only with *structured problems*, where the elements of the problem and their relations are well understood by the decision-maker. By contrast, *non-programmed decisions* cannot make use of standard procedures, either because similar problems have not arisen in the past, or because of their imprecise and complex nature. Such problems are called *unstructured*, and neither the elements nor their relationships are understood by the problem solver. Simon readily accepts that real-life decisions cannot be neatly classified into these two categories, which really represent the two extremes of a continuum. As a corollary, many problems must be seen as *semi-structured*; that is, some elements and relations are understood by the problem solver, but others are not. Again, we can see the importance of bounded rationality in this classification. Programmed decisions can be made fairly easily because previous solutions, as well as the cues indicating their appropriate application, have been stored in long-term memory; this luxury is not possible with non-programmed decisions, and then a relatively slow and serial problem solving behavior must be used.

The thesis put forward by Simon is that bounded rationality is inevitable due to the limits imposed on human thinking by limits of attention, learning, and memory capacity. As seen above, a major way to alleviate these limits is learning, that is to use long-term memory as an extension of short-term memory. In extreme cases, learning will lead to high levels of expertise. Rather than storing a handful of items in short-term memory, one can store a large amount of information by using chunks and schemata. Rather than carrying attention to a single action, one can carry out several actions in parallel without using much attention, assuming that each action has been learned to the point that it has become automatized.

Simon has developed the link between learning and bounded rationality in the *chunking theory*, which assumes that learning involves the acquisition of large numbers of perceptual chunks <<Link to Chunking mechanisms and learning>> and actions linked to these chunks. This *large* reservoir of prior experience

injects flexibility and efficiency into a system that would otherwise be rigid. Simon estimated that a typical human expert <<Link to Development of expertise>> will have acquired between 10,000 and 100,000 chunks over a training period of 10 years. This knowledge, together with the presence of efficient heuristics, reduces the need for search, as previous searches have essentially been compiled in long-term memory. Efficient access and immediate response give rise to the intuitive ease with which domain experts rapidly diagnose and accurately respond to problems in their area of expertise.

Important Scientific Research and Open Questions

The concept of bounded rationality is amply supported by empirical evidence on probabilistic reasoning, logical reasoning, and game playing, to cite just a few domains outside organizational behavior. While fairly well accepted in organizational science and psychology, Simon's ideas are still disputed in economics, where (neo-)classical theories are still dominant. It should also be recognized that the idea of rationality, in the sense of maximizing expected utility, can still be found in fields such as psychology, most notably in Anderson's framework of rational analysis (Anderson, 1990).

Recent research on bounded rationality has followed several avenues. The extent to which heuristics and algorithms help alleviate the constraints imposed by humans' cognitive limitations has received particular attention (e.g. Gigerenzer, 2000). A counter-intuitive result from this line of research is that simple heuristics sometimes lead to decisions that are better than optimal decision algorithms, such as linear regression. Another line of research has led to the development of formal models – either mathematical or computational – to account for the way decisions are made (e.g. Rubinstein, 1998).

As noted above, learning is one of the two main ways humans have expanded their bounded rationality, and there is also important research on this topic. Two lines were directly started by Simon himself. The first line, on expertise, tries to investigate the limits of human cognition and to understand the training mechanisms that allow experts to push these limits much further than novices. Interestingly, while the assumption of full rationality is at variance with research documenting the development of expertise – differences in expertise are meaningless in this framework, because even 'novices' should have unlimited computational resources and access to information – bounded rationality offers a natural theoretical account to understand expertise and its development (Campitelli & Gobet, in press). The second line has developed computer models of the way humans acquire knowledge, and has focused on the acquisition of perceptual chunks. One of Simon's key insights is that acquired knowledge enables humans to take excellent and rapid decisions by intuition, thus side-stepping the limits of their limited rationality. The computer model CHREST <<Link to Learning in the CHREST cognitive architecture>> (Gobet et al., 2001) contains mechanisms explaining how this can happen.

Cross-References

→ Learning in the CHREST cognitive architecture
→ Schema
→ Decision making and learning
→ Chunking mechanisms and learning
→ Development of expertise

References

Anderson, J. R. (1990). *The adaptative character of thought*. Hillsdale, NJ: Erlbaum.

Campitelli, G., & Gobet, F. (in press). Herbert Simon's decision-making approach: Investigation of cognitive processes in experts. *Review of General Psychology*.

Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. New York, NY: Oxford University Press.

Rubinstein, A. (1998). *Modeling bounded rationality*. Cambridge, MA: MIT Press.

Simon, H. A. (1947). *Administrative Behavior*. New York, NY: Macmillan.

Simon, H. A. (1997). *An empirically based microeconomics*. Cambridge, UK: Cambridge University Press.

Definitions

Bounded rationality: The notion that, due to cognitive limitations in knowledge and computational capacity, decision makers cannot aim for optimal choices, but have to be content with good-enough solutions – that is, they have to "satisfice."

Decision: A choice between two or more alternatives. A central question in the social sciences, and in particular economics, is whether humans make decisions that are optimal (global rationality) or whether they have to aim only for decisions that are "good-enough" (bounded rationality).

Global rationality: The assumption, made by social-science approaches such as neo-classical economics, that humans can make optimal decisions.

Procedural rationality: The notion that rationality should be studied by focusing on the decision processes and on the decision maker. Sometimes used as a synonym for bounded rationality.

Satisficing: The notion that a decision maker stops their search and selects a solution that satisfies some criteria, rather than continuing to search for an optimal solution.

Substantive rationality: The notion that rationality should be studied by focusing on an abstract analysis of the situation rather than on decision makers and the processes they use. Sometimes used as a synonym for global rationality.