

Decision Making: Nonrational Theories

Gerd Gigerenzer

The term “nonrational” denotes a heterogeneous class of theories of decision making designed to overcome problems with traditional “rational” theories. Nonrational theories have been denoted by various terms, including models of bounded rationality, procedural rationality, and satisficing. Although there is as yet no agreed-upon definition of “nonrational,” rational and nonrational theories typically differ on several dimensions, discussed below. The term “decision making” is used broadly here to include preference, inference, classification, and judgment, whether conscious or unconscious.

Nonrational theories of decision making should not be confused with theories of irrational decision making. The label “nonrational” signifies a type of theory, not a type of outcome. In other words, the fact that nonrational theories postulate agents with emotions, limited knowledge, and little time—rather than postulating omniscient “rational” beings—need not imply that such agents fare badly in the real world.

1. Historical Background

A few historical remarks on rational theories help to set the stage for nonrational theories. In the mid-seventeenth century, the calculus of probability replaced the ideals of certain knowledge and demonstrative proof (as in mathematics and religion) with a more modest vision of reasonableness. What may be called the first rational theory of decision making, the maximization of expected value, emerged at this time. According to the theory, an option’s expected value is the sum of the product of the probability and the value of each of its consequences; a rational decision maker chooses the option with the highest expected value.

The notion of defining an option’s reasonableness in terms of its expected value soon ran into problems, because in some situations (e.g., the St. Petersburg problem), its prescriptions conflicted with educated intuition. Mathematician Daniel Bernoulli therefore proposed to replace the concept of expected value with that of expected utility. For instance, the utility of a monetary gain (say, of \$1,000) can be defined as a logarithmic function of its dollar value and the agent’s current wealth, assuming that the utility of an additional dollar diminishes as the value of the gain and current wealth increase.

The fact that rational decision making can be defined in more than one way—for example, as maximization of expected value or expected utility—has been interpreted both as the weakness and the strength of the program. This ambiguity was one of the reasons why, by 1840, most mathematicians had given up attempting to define a calculus of reasonableness (Daston, 1988). With a few exceptions, rational theories of decision making largely disappeared until their revival in the 1950s and 1960s. Only then did the major species of rational theories, the maximization of subjective expected utility and Bayesianism, become influential in the social and behavioral

sciences. At about the same time, some psychologists and economists—most notably Nobel laureates Herbert Simon and Reinhard Selten—criticized the assumptions about the human decision maker in modern rational theories as empirically unfounded and psychologically unrealistic, and called for alternative theories.

2. Optimizing vs. Nonoptimizing Theories

Rational theories rest on the ideal of optimization; nonrational theories do not. Optimization means the calculation of the maximum (or minimum) of some variable across a number of alternatives or values. For instance, according to a rational theory known as subjective expected utility (SEU) theory, an agent should choose between alternatives (e.g., houses, spouses) by determining all possible consequences of selecting each alternative, estimating the subjective probability and the utility of each consequence, multiplying the probability by the utility, and summing the resulting terms to obtain that alternative's subjective expected utility. Once this computation has been performed for each alternative, the agent chooses the alternative with the highest expected utility. This "subjective" interpretation of SEU has been used to instruct people in making rational choices, but was criticized by decision theorists who argue that preferences are not derived from utilities, but utilities from preferences. In this "behavioristic" interpretation, no claims are made about the existence of utilities in human minds; SEU is only an attempt to describe the choice. People choose *as if* they are maximizing SEU (see section 3).

Nonrational theories dispense with the ideal of optimization. For instance, Simon (e.g., 1956, 1982) proposed a nonrational theory known as *satisficing*, in which an agent is characterized by an aspiration level and chooses the first alternative that meets or exceeds this aspiration level. The aspiration level (e.g., characterization of what would constitute a "good-enough" house) allows the agent to make a decision without evaluating all the alternatives.

There are several motives for abandoning the ideal of optimization. First, in many real-world situations, no optimizing strategy is known. Even in a game such as chess, which has only a few stable, well-defined rules, no optimal strategy exists that can be computed by a human or a machine. Second, even when an optimizing strategy exists, it may demand unrealistic amounts of knowledge about alternatives and consequences, particularly when the problem is novel and time is scarce. Acquiring the requisite knowledge can conflict with goals such as making a decision quickly; in situations of immediate danger, attempting to optimize can even be deadly. In social and political situations, making a decision can be more important than searching for the best option. Third, strategies that do not involve optimization can sometimes outperform strategies that attempt to optimize. In other words, the concept of an optimizing strategy needs to be distinguished from the concept of an optimal outcome. In the real world, there is no guarantee that optimization will result in the optimal outcome. One reason is that optimization models are built on simplifying assumptions that may or may not actually hold. An example of a nonoptimizing strategy that performs well in the repeated prisoner's dilemma is "tit for tat," a simple heuristic that cooperates on the first move and thereafter relies on imitation, cooperating if the partner cooperated and defecting if the partner defected on the previous move. In the finitely repeated prisoner's dilemma, two tit-for-tat players can make more money than two rational players (who reason by "backward induction" and therefore always defect), although tit for tat only requires remembering the partner's last move and does not involve optimization.

3. Normative vs. Descriptive Theories

Nonrational theories are descriptive, whereas rational theories are normative—this common distinction is half-true. Indeed, nonrational theories are concerned with psychological plausibility, that is, the capacities and limitations of actual humans, whereas rational theories have little concern for descriptive validity and tend to assume omniscience. But nonrational theories have sometimes been interpreted as normative as well. For instance, if an optimization strategy is nonexistent, unknown, or dangerous to perform because it would slow decision making, a simple heuristic—such as copying the behavior of others—can be the best decision-making strategy.

Rational theories typically do not assume that agents actually perform optimization or have the knowledge needed to do so. Their purpose is not to describe the reasoning process, but to answer a normative question: what would be the best strategy for an omniscient being to adopt? In economics, psychology, and animal behavior, however, the answer to this question is sometimes used to predict actual behavior. In this way, a rational theory can be descriptive of behavioral outcomes yet mute about underlying processes. For instance, optimal foraging theory assumes that animals select and shift between food patches *as if* they had perfect knowledge about the distribution of food, competitors, and other relevant information, without claiming that real animals have this knowledge or perform optimizing computations. Instead, it is assumed that animal behavior has evolved to be close to optimal in specific environments. The question of what proximal mechanisms produce this behavior is a different one. These mechanisms may be heuristics, habits, or forms of social imitation that are the topic of theories of nonrational decision making.

To summarize, nonrational theories aim to describe both the process and the outcome of decision making. In certain situations, they can be seen as characterizing the best an organism with limited time and knowledge can do. Rational theories are primarily normative, although often for omniscient rather than real beings. They are often seen as descriptive in the sense of predicting behavior, but not as models of underlying processes.

4. Search vs. Omniscience

Search is a central part of nonrational theories. In many rational theories, in contrast, search plays no role; it is instead assumed that all relevant information is already available to the agent (hence the agent's omniscience). Actual humans, however, have to search for information, either in memory or in outside sources, such as on the Internet or in encyclopedias. Information search can cost time and money.

One class of rational theories, known as “optimization under constraints,” models limited search but retains the ideal of optimization. These theories posit an optimal stopping rule that requires the organism to stop search when the costs of further search exceed its benefits. This rule has been criticized because it can lead to infinite regress: To compute the optimal trade-off between the costs and benefits of search itself carries costs, which must be factored in to a “meta” analysis of the costs and benefits of computing the costs and benefits of search, and so on. A second point of criticism is that accurate estimates of benefits and costs—such as opportunity costs, incurred by foregoing the benefits of activities other than continuing search—can be troublesome to obtain and may even demand knowledge close to omniscience. Therefore, optimization under constraints can lead to models that are descriptively even more unrealistic than rational theories that ignore search.

Nonrational theories model search with simple stopping rules rather than optimal stopping rules. Search can concern either of two kinds of information, *alternatives* (such as houses and

spouses) or *cues* (such as reasons to choose a given house). Two different classes of nonrational theories deal with these types of search: aspiration level theories with the search for alternatives and fast and frugal heuristics with the search for cues.

4.1 Aspiration Level Theories

Aspiration level theories assume that an agent has an aspiration level, which is either a value on a goal variable (e.g., profit or market share) or, in the case of multiple goals, a vector of goal values that is satisfactory to the agent. When choosing among a large (possibly even infinite) set of alternatives, agents search until they find the first alternative that meets or surpasses their aspiration level, at which point search stops and that alternative is chosen. For instance, agents might set a lower limit on the price at which they would be willing to sell their shares in a company (the aspiration level). In this satisfying model, the agent makes no attempt to calculate an optimal stopping point, in this case, the best day on which to sell. The aspiration level need not be fixed, but can be dynamically adjusted to feedback. For instance, if the investor observes that the share price is monotonically increasing rather than fluctuating over time, they might conclude that there is some stable trend and adjust the limit upward. Thus, aspiration level theories model decision making as a dynamic process in which alternatives are encountered sequentially and aspiration levels stop search. The challenge is to understand where aspiration levels come from in the first place (Simon, 1982; Selten, 1998).

4.2 Fast and Frugal Heuristics

In a different class of problems, the set of alternatives is given (i.e., need not be searched for), and the agent needs to search for cues that indicate which alternative to choose. For instance, an employer might want to decide which of three job applicants to hire or a bettor to predict which of two soccer teams will win a game. *Fast and frugal heuristics* employ simple stopping rules to make such inferences with little computation (“fast”) and information (“frugally”). For instance, the “take the best” heuristic bases its inference solely on the best cue on which the alternatives differ and ignores the rest. Such “one-reason” decision making allows agents to make decisions quickly and, counterintuitively, often as or more accurately than the “optimal” linear model (multiple regression), which looks up and integrates all available cues. Other types of fast and frugal heuristics include ignorance-based decision making (see below) and elimination heuristics. Thus, fast and frugal heuristics model decision making as a dynamic process in which cues or reasons are sequentially searched for—in memory or in the outside world—and simple stopping and decision rules determine inferences. The challenge here is to understand what the class of heuristics is, how a heuristic is selected, and in which environments it is successful (Gigerenzer et al., 1999).

5. Ecological Rationality vs. Internal Consistency

A classical criterion for rational choice is internal consistency or coherence. Numerous rules of consistency have been formulated, for instance, transitivity and additivity of probabilities. These rules, which are the building blocks of rational theories, have been used to test the limits of human rationality, beginning with the work of Jean Piaget and Bärbel Inhelder. Nonrational theories, in

contrast, place less weight on internal consistency; for instance, some fast and frugal heuristics can violate transitivity. Instead, nonrational theories emphasize performance in the external world, both physical and social. Measures of external performance include the accuracy, speed, frugality, cost, transparency, and justifiability of decision making. Note that internal consistency does not guarantee that any of these external criteria are met. For instance, the statement “there is a 0.01 probability that cigarette smoking causes lung cancer and a 0.99 probability that it does not” is internally consistent in that the probabilities add up to 1, but according to relevant research, it is not accurate.

How can heuristics be simple and accurate at the same time? Two major answers have been proposed: They can exploit the structure of environments, and simplicity can entail robustness.

5.1 *Structure of Environments*

The term “ecological rationality” refers to the match between a heuristic and the structure of the information in a particular environment. The more closely a heuristic reflects important aspects of this structure, the more likely it is to succeed in that environment. Simple heuristics can succeed by exploiting the structure of information in an environment. In other words, the environment can do part of the work for the heuristic. For instance, consider the problem of predicting which of two soccer teams will win a game, which of two cities is larger, or which of two colleges provides the better education. Assume a fairly ignorant agent who has heard of only one of the two teams, cities, or colleges. He can use the “recognition” heuristic, which infers that the recognized object will win the game, have the larger population, or provide the better education (Gigerenzer et al., 1999). Such ignorance-based decision making works well in environments where ignorance (e.g., lack of name recognition) is not random but systematic, as in competitive environments where the sequence in which the names of objects are first encountered is correlated with their performance, power, or size. The structure of such environments does part of the work in the sense that it allows the recognition heuristic to glean information from ignorance. If the correlation between recognition and the criterion is sufficiently large, a counterintuitive result is observed: Less knowledge leads to more accurate predictions than more knowledge, because people who recognize both alternatives cannot use the recognition heuristic.

Cricket, baseball, and other sports in which players need to catch a ball provide a second illustration of how heuristics can exploit the structure of environments. Suppose that we want to build a robot that can catch a ball—a thought experiment because none exists so far. One approach to building robots in artificial intelligence is to endow them with a complete representation of the environment and a supercomputer that can draw inferences from this information. Taking this approach, which is consistent with rational theories, one would feed the robot parabolic functions describing all the trajectories that the ball might follow, and equip it with instruments for measuring the ball’s initial velocity and angle to select the right parabola. Further instruments would be needed to measure the speed and direction of the wind at each point of the ball’s flight, as well as myriad other factors, such as spin, in order to calculate the ball’s deviation from the parabolic course. These measurements would be analyzed by a computer that would then compute where the ball will land. Professional athletes, in contrast to this hypothetical robot, seem to use a simple heuristic for catching (McLeod & Dienes, 1996). They start running with their eyes fixed on the ball and adjust their speed to keep their angle of gaze constant (or within a certain range). Using this heuristic, a robot can intercept the ball without searching for the information that a “rational” robot requires. It pays attention to only one variable—the angle of gaze, which does all the work. The “heuristic” robot does not calculate where the ball will land, but it will be there when it does.

These examples illustrate how heuristics can succeed by exploiting structures of environments. These heuristics are “domain-specific” rather than “domain-general,” that is, they work in a class of environments in which they are ecologically rational. Heuristics do not provide a universal rational calculus, but a set of domain-specific mechanisms similar to the parts of a Swiss army knife, and have been referred to collectively as the “adaptive toolbox” (Gigerenzer & Selten, 2001).

5.2 Robustness

A second reason why a simple heuristic can make accurate predictions is robustness. To understand this point, it is necessary to distinguish between data fitting (i.e., determining the best-fitting parameter values for a model given a specific body of data) and prediction (i.e., using these parameter values to predict new data). For data fitting, it generally holds that the more parameters one uses in a model, the better the fit; for prediction; however, there can be a point where less is more (Foster & Sober, 1994). For instance, if one records the air temperature on all 365 days in a year, one can fit the resulting jagged curve increasingly well as one adds more free parameters to the model. However, if one wants to predict air temperature during the coming year, the model that best fits the past data may result in less accurate predictions than a simpler model with fewer parameters and a worse fit. Similarly, one can fit one ball’s trajectory through the air to arbitrary degrees of precision, but this may be of little help in predicting the next ball’s trajectory.

More generally, in noisy environments only part of the available information generalizes to the future. The art of building a good model is to find this part and to ignore the rest. The more noise in the environment, the more that models with many free parameters tend to “overfit,” that is, to reflect the noise. Overfitting can become a problem when overly powerful mathematical models, such as neural networks with numerous hidden units and multiple regression with many predictors, are used to fit and then predict behavioral data (Geman et al., 1992). Simplicity can reduce overfitting and thereby produce robust decision strategies. An alternative route to robustness is to use the computational power of modern computers to search through large numbers of models to find one that is robust by a given criteria. But one often does not have the time and knowledge to proceed this way.

Heuristics can be fast, frugal, and accurate by exploiting the structure of information in environments and by being robust. But these are not the only reasons why strategies that work with limited knowledge and time can be successful. For instance, there is evidence that language learning fails to occur in fully formed neural networks, but successfully develops in networks that begin with limited working memory and gradually mature. This finding has been interpreted to mean that young children’s memory limitations may be an important precondition for rapid language acquisition (Elman, 1993). Thus, cognitive limitations might not always be regrettable deficits, but actually enable learning by starting small.

6. Emotions, Imitation, and Social Norms

Like rational theories, most nonrational theories rely on cognitive building blocks, such as aspiration levels, search heuristics, and stopping rules. However, *Homo sapiens* is not only the most intelligent, but also the most emotional and social species—one of the very few in which unrelated members of the same species cooperate. Theories of decision making have often neglected

emotions, and sometimes even cast them as the very opposite of rationality. However, emotions can aid decisions making just as fast and frugal cognitive heuristics do. For instance, falling in love can be seen as a powerful stopping rule that ends search for a partner and strengthens commitment to the loved one. This emotion guarantees commitment more effectively than would a nonemotional mind that tried to optimize partner quality and dropped the current partner each time a more promising one came along. Similarly, feelings of parental love, triggered by the presence or smile of one's infant, can prevent cost-benefit computations, so that the question of whether it is worthwhile to endure all the sleepless nights and other challenges of baby care simply never arises. Instead, love keeps parents focused on the adaptive task of protecting and providing for their offspring. To take another emotion as an example, disgust can limit the choice set of potential foods and help one to avoid becoming ill from spoiled food. Finally, emotions such as fear and anger can speed decision making to the point that there is no decision to be made.

Like emotions, social norms and social imitation can function as decision-making guidelines that keep individual learning and information search to a minimum (Gigerenzer & Selten, 2001). Social heuristics such as "eat what other conspecifics eat" or "prefer mates preferred by others" can guide behavior without much information gathering and bring benefits such as a reduced likelihood of food poisoning and social disapproval. These forms of social rationality can be found throughout the animal world. For instance, female guppies prefer the mates that other females prefer, and such social imitation can even reverse their prior preferences for one male over another. In humans, media-fueled mate preferences, and occasionally even academic hiring, follow similar heuristics, such as "If they want him, then we want him!" Custom, not optimization, governs much of life, even in the economic and intellectual worlds.

Social systems foster not only individual but distributed intelligence. That is, by cooperating with one another, myopic individuals can exhibit collective rationality. Communities of social insects are one example of such intelligent "superorganisms," as is division of labor in human industry and politics. The intelligence of a superorganism can be seen as the emergent property of a few adaptive heuristics of its members. Honey bees, for instance, make intelligent collective decisions about where to build a new hive that seem to emerge from individual bees' application of a few simple, well-adapted heuristics. Complex phenomena need not to be modeled in terms of complex knowledge and computation.

7. Summary

Nonrational theories take account of what we know about humans' and other species' capacities rather than assuming unlimited knowledge, memory, time, and other resources. They model heuristics—cognitive, emotional, and social—that exploit the structure of information in real environments. Nonrational theories conflict with the ideal of *Homo economicus* and other visions of humans developed in the image of an omniscient God, but provide us with a more realistic picture of decision making when knowledge is scarce, deadlines are rapidly approaching, and the future is hard to predict.

See also: Bayesian Theory: History of Applications; Bounded and Costly Rationality; Bounded Rationality; Decision and Choice: Random Utility Models of Choice and Response Time; Decision Biases, Cognitive Psychology of; Decision Theory: Bayesian; Heuristics for Decision and Choice; Luce's Choice Axiom; Rational Choice Explanation: Philosophical Aspects; Subjective Probability Judgments; Utility and Subjective Probability: Contemporary Theories; Utility and Subjective Probability: Empirical Studies.

Bibliography

- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
- Daston, L. (1988). *Classical probability in the Enlightenment*. Princeton, NJ: Princeton University Press.
- Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, *34*, 571–582.
- Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*, 71–99.
- Foster, M., & Sober, E. (1994). How to tell when simpler, more unified, or less ad hoc theories will provide more accurate predictions. *British Journal of Philosophical Science*, *45*, 1–35.
- Geman, S., Bienenstock, E., & Doursat, E. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, *4*, 1–58.
- Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. Cambridge, MA: MIT Press.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- McLeod, P., & Dienes, Z. (1996). Do fielders know where to go to catch the ball or only how to get there? *Journal of Experimental Psychology: Human Perception and Performance*, *22*, 531–543.
- Piaget, J., & Inhelder, B. (1951/1975). *The origin of the idea of chance in children*. New York: Norton.
- Rivest, R. J. (1987). Learning decision lists. *Machine Learning*, *2*, 229–246.
- Selten, R. (1998) Aspiration adaptation theory. *Journal of Mathematical Psychology*, *42*, 191–214.
- Selten, R., & Buchta, J. (1999). Experimental scaled bid first price auctions with directly observed bid functions. In D. Budescu, I. Erev, & R. Zwick (Eds.), *Games and human behavior: Essays in honor of Amnon Rapoport* (pp. 79–102). Mahwah, NJ: Erlbaum.
- Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review*, *63*, 129–138.
- Simon, H. A. (1982). *Models of bounded rationality*. Cambridge, MA: MIT Press.
- Wimsatt, W. C. (in press). *Re-engineering philosophy for limited beings: Piecewise approximations to reality*. Cambridge, MA: Harvard University Press.