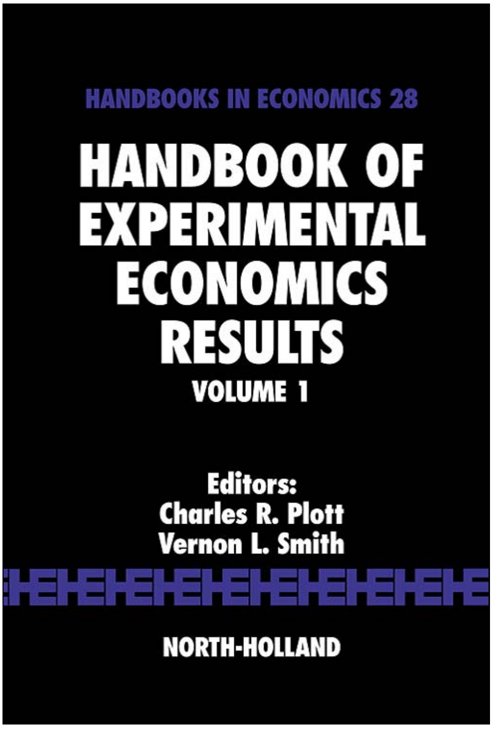

Gerd Gigerenzer and Peter M. Todd, Rationality the Fast and Frugal Way: Introduction. In: C.R. Plott and V.L. Smitt, editors, Handbook of Experimental Economics Results, Volume 1. Amsterdam: North-Holland, 2008, p. 976.
ISBN: 978-0-444-82642-8

Chapter 105

RATIONALITY THE FAST AND FRUGAL WAY: INTRODUCTION

GERD GIGERENZER

PETER M. TODD

Max Planck Institute for Human Development, Berlin, Germany

What is bounded rationality? The neoclassical answer is optimization under constraints such as information costs (e.g., Sargent, 1993). For behavioral economists, however, bounded rationality is the study of cognitive illusions and decision anomalies (e.g., Camerer, 1998; Rabin, 1998). These two interpretations make an odd couple, one promoting rationality, the other irrationality. Behavioral economists report that human behavior deviates from optimization models and reject these as descriptively inadequate. Proponents of optimization respond with Milton Friedman's (1953) "as if" defense: Our models may incorrectly depict people as if they had the statistical sophistication of econometricians, but this does not matter, because our "as if" models are not intended to portray the process of decision making, but only to predict its outcome. One might think that behavioral economists focus on process, but that is not so either. Prospect theory (Kahneman and Tversky, 1979), for instance, predicts outcomes, not processes. Its equations are not meant to describe the steps of decision making, nor is there experimental evidence that people would actually go through these calculations. The same holds for other models that leave the Bernoullian framework of probability-times-utility untouched but tinker with the functions or add a single variable such as "regret" or "disappointment." If one is interested in how decisions are actually made, and in using this knowledge to design prescriptive models for improving decision making, one must go back to the blackboard and risk a more radical step.

This step is inherent in a third interpretation of bounded rationality that focuses on process: the study of the fast and frugal heuristics people use to make decisions with limited time and knowledge (Gigerenzer and Selten, 2001a; Gigerenzer, Todd, and the ABC Research Group, 1999). This third interpretation is an elaboration of Simon's (1955, 1990) analogy between bounded rationality and a pair of scissors: one blade is cognition and its limitations, the other the environment and its structure. To focus on bounded rationality as the study of heuristics that can exploit structures of their environments, we use the term "ecological rationality." If one focuses only on cognitive limitations, one can hardly understand why cognition works as well as it does – just as looking at one blade alone does not explain how scissors cut. Simon's use of the term "cognitive limitations" (as opposed to neutral terms such as "cognitive heuristics") is unfortunate because it has been taken to imply poor performance. However, equating limits with failure, and lack of limits with success, may underlie a deep misunderstanding about the consequences of omniscience. An unlimited memory, for instance, could

Handbook of Experimental Economics Results, Volume 1

DOI: 10.1016/S1574-0722(07)00105-9

be disastrous: The sheer mass of details stored could critically slow down and inhibit the retrieval of the few important experiences when they are needed. And too much information would impair the mind's ability to abstract, infer, and learn. A human brain or an artificial neural network that starts with an overabundant memory capacity may never be able to learn a language, whereas having a working memory that starts small (and then grows) can act as a filter of environmental input, enabling learning of important linguistic scaffolding first (Elman, 1993; Hertwig and Todd, 2003). Less can be more.

1. Heuristics

The term "heuristic" is of Greek origin, meaning "serving to find out or discover." Heuristics are a useful instrument in science as well as in everyday decision making (Payne, Bettman, and Johnson, 1993; Polya, 1954). Einstein (1905 taken from Holton, 1988) used the term "heuristic" to label an idea that he considered incomplete, given the limits of our knowledge, but useful. Behavioral economists, in contrast, have been using the term "heuristics" to account for errors of judgment and biased behavior that should be avoided. The positive power of heuristics has been replaced by a bad reputation. Moreover – or perhaps as a consequence – there are rather few concrete models of heuristics, and instead a preponderance of vague labels such as "availability" that can be applied post hoc to explain everything and nothing (see the exchange between Gigerenzer, 1996 and Kahneman and Tversky, 1996).

In this chapter, we give an introduction to the study of fast and frugal heuristics as specific models of decision mechanisms; the chapters that follow focus on particular heuristics. Fast and frugal heuristics do not try to optimize, that is, to compute the maximum or minimum of some function, nor, for the most part, do they calculate probabilities or utilities. Rather, they rely on simple building blocks for searching for information, stopping search, and finally making a decision. The goal of studying such heuristics is both descriptive and prescriptive. First, understanding how people make decisions enables us to predict when they will succeed and fail, typically more effectively than an outcome-oriented "as if" model would allow. Second, knowing which heuristics work in which environments allows us to design appropriate heuristic decision procedures for use by institutions, individuals, and artificial systems, and to teach people how to decide in an uncertain world when time is pressing and resources are scarce.

2. A Fast and Frugal Heuristic

We illustrate the mesh between the descriptive and prescriptive goals with an example from sports. Imagine a company that wants to design a robot that can catch balls, as in cricket and baseball. For the sake of simplicity, we will only consider the case where a ball comes in high, behind or in front of a player. (This is a thought experiment, as no such robot exists yet; see Gigerenzer and Selten, 2001b.)

One team of engineers, taking an optimizing approach, attempts to provide the robot with expensive computing and sensory equipment and a full representation of its environment (omniscience), including knowledge of possible parabolic flight paths and the ball's initial distance, velocity, and projection angle. But in a real game, due to air resistance, wind, and spin, balls do not fly in parabolas. Hence, the robot would need more instruments that can measure the speed and direction of the wind at each point of the ball's flight in order to compute the resulting path and ultimately the spot where the ball will land.

A second team, the cognitive illusion team, contends that the optimization equation is descriptively wrong. Real players are bounded in their rationality, and so make systematic errors. They show that when experienced players are asked to predict where a fly ball will land, they consistently underestimate the distance to the landing spot. They call this the "optimistic bias" – because underestimating the distance suggests to players that they might actually get the ball even when they cannot. The optimization team responds that they will nevertheless maintain their "as if" model; a model that can approximately predict the point to which players run is better than no model. Furthermore, they argue, even if the story about the optimistic bias were true, it would not help to understand how actual players catch a ball, nor how to build the robot.

A third group of engineers, the heuristics team, agrees that humans may not be able to compute the point where the ball will land. However, they point out, that may not be the interesting question. Players, after all, typically can catch high fly balls, so the proposed optimism bias is not holding them back. Instead, the question is: how do players catch a ball, if they do not perform the measurements that the optimization team proposes? Experiments and observation have shown that experienced players use a fast and frugal heuristic (McLeod and Dienes, 1996). When a ball comes in high, the player visually fixates the ball and starts running. The heuristic is to adjust the running speed so that the angle of gaze (between the ball and the horizon) remains roughly constant. In our thought experiment, a robot that uses this heuristic does not need to measure wind, air resistance, spin, or any other causal variables. It can get away with ignoring this information. All the relevant information is contained in one variable: the gaze angle. The gaze heuristic thus uses one-reason decision making. Note that this robot is not able to compute the point at which the ball will land. But it will be there when the ball comes low enough to catch. What looks like a serious mental flaw in need of de-biasing turns out to be irrelevant for good ball catching.

How is such a simple solution possible? The gaze heuristic makes use of an environmental feature that the player can easily exploit. In place of the complicated true trajectory of the ball's flight – which the optimization team was trying to work out – the gaze heuristic uses (and, through the action it dictates, creates) a simple relation between the player's eye and the ball. This solution is an instance of ecological rationality; that is, exploiting a match between cognition and its environment.

This thought experiment illustrates that a good model of the decision process (here the gaze heuristic) can supersede "as if" optimization models, both descriptively and prescriptively. First, the heuristic can predict behavior that the "as if" model cannot, such as that the player will catch the ball while running because he has to move to keep the angle of gaze constant. This is an experimentally testable prediction, and in fact, players do not run to a spot and wait there – they catch the ball while running. Second, a good process model can predict what the person relying on a heuristic cannot do, such as computing the point where the ball will land. This will help behavioral economists to predict what decision failures can occur and why. Third, a heuristic can aid in achieving prescriptive goals. For instance, the gaze heuristic can be used to make a robot good at ball catching, or to instruct inexperienced players. An optimizing "as if" model, however, may lead to computational explosion that makes it impossible to implement in any hardware, whether human or computer.

3. The Adaptive Toolbox

The vision of rationality embodied in the first team of robot builders matches that held by many economists: rationality as optimization, with or without constraints, assuming that economic agents act as if they are omniscient (at least approximately) and have the ability to make sophisticated statistical calculations. The Bayesian approach and subjective expected utility maximization are examples of this vision. In contrast, the vision of ecological (bounded) rationality takes into account the realistic psychological abilities of agents and the structures of the environments they face. Instead of relying on one all-powerful decision mechanism, ecological rationality arises through the use of a collection of fast and frugal heuristics in the mind's adaptive toolbox. The gaze heuristic is one such tool in this box. Like hammers and wrenches, heuristics are designed for specific classes of problems. The gaze heuristic, for instance, works for a class of problems that involve the interception of moving objects. If you learn to fly an airplane, you will be taught a version of it for *avoiding* interception of another moving plane. In other domains, such as choosing what to eat, this heuristic would not be sensible to use.

We call the heuristics "fast" because they process information in a relatively simple way, and "frugal" because they search for little information. To study these heuristics, they are formulated as computational models, making them amenable to mathematical analysis and simulation as well as experimentation. (This is one major difference from the heuristics-and-biases program – e.g., Kahneman and Tversky, 1996 – which did not specify precise mechanisms.) Each heuristic is composed of building blocks, or heuristic principles, that serve three functions: They give search a direction, stop search, and make a decision.

3.1. Heuristic Principles for Guiding Search

Alternatives and cues must be sought in a particular order. For instance, search for cues can be simply random or in order of cue validity.

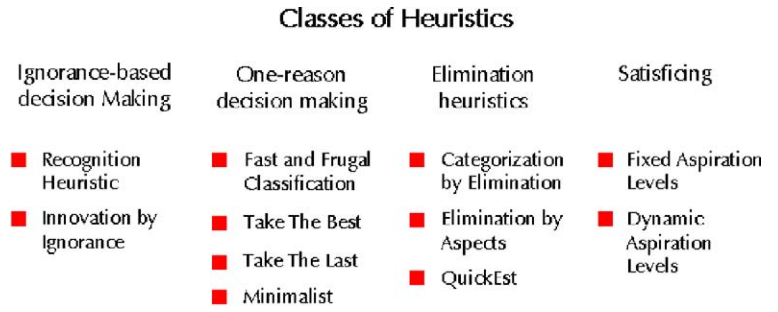

Figure 1. Four classes of simple heuristics distinguished by the problem settings in which they are applied, along with examples of each that can be constructed with different combinations of building blocks for searching, stopping, and deciding.

3.2. Heuristic Principles for Stopping Search

Search for alternatives or cues must be stopped at some point. Fast and frugal heuristics employ stopping rules that do not try to compute an optimal cost-benefit trade-off as in optimization under constraints. Rather, heuristic principles for stopping involve simple criteria that are easily ascertained, such as halting information search as soon as the first cue or reason that favors one decision alternative is found.

3.3. Heuristic Principles for Decision Making

Once search has been stopped, a decision or inference must be made. Many models of judgment and decision making ignore search and stopping rules and focus exclusively on decision: Are predictor values combined in a Bayesian way, linearly as in multiple regression, or in some other fashion? Instead, we focus on simple principles for decisions (such as one-reason decision making) that avoid expensive computations and extensive knowledge by working hand in hand with equally simple search and stopping rules.

Figure 1 shows four (non-exhaustive) classes of heuristics in the adaptive toolbox that can be built from these building blocks. Ignorance-based decision making refers to heuristics that can exploit the fact that one does not know everything. For instance, consumer behavior is guided by brand name recognition, a necessary precondition for reputation. When consumers choose between two brands, one they recognize by name, the other not, they tend to select the recognized one. This recognition heuristic works in environments where name recognition is correlated with product quality, but can be exploited by firms who put their money into promoting name recognition through advertisement rather than increasing the quality of the product. The next two chapters introduce the recognition heuristic, the less-is-more effect, and studies in which the heuristic guides investment decisions in the stock market. Heuristics falling into the second class, one-reason decision making, are introduced in Chapter 108. The heuristics in this class either pay attention only to one reason or cue, such as the gaze heuristic, or they search through several reasons but at the end only use one to reach their decision. The third class of elimination heuristics contains mechanisms that start with a set of alternatives and go through a sequence of cues to successively eliminate alternatives until only one remains (Berretty, Todd, and Martignon, 1999; Tversky, 1972). Finally, satisficing heuristics in the fourth class go through alternatives (rather than cues) sequentially – for example in looking for a new home – and choose the first one that satisfies an aspiration level, which may change with time spent on search (Selten, 2001; Simon, 1990). The heuristics in these four classes are described in more detail in Gigerenzer and Selten (2001a), Gigerenzer, Czerlinski, and Martignon (1999), and Todd and Gigerenzer (2000).

4. Emergency Room Decisions

In the following chapters we present examples of various heuristics from the adaptive toolbox. Here we begin with an example of how a heuristic for a particular, very important application can be constructed from the three types of building blocks described above, illustrating the prescriptive use of simple decision mechanisms.

A man is rushed to the hospital with serious chest pains. The doctors suspect acute ischemic heart disease and need to make a decision – quickly: Should the patient be assigned to the coronary care unit or to a regular nursing bed for monitoring? How do doctors make such a decision, and how should they? In a study conducted at a Michigan hospital, doctors relied on their intuition and sent some 90% of the patients to the coronary care unit (Green and Mehr, 1997). This defensive decision making led to unnecessary costs (too many people in the expensive coronary care unit), decreased the quality of care provided (the unit was overcrowded), and became a health risk for patients who should not be in the unit (it is one of the most dangerous places in a hospital due to the risk of secondary infections, which can be fatal). Other researchers tried to solve this overcrowding problem by training physicians to use the Heart Disease Predictive Instrument (Pozen et al., 1984). This decision support system consists of a chart with some 50 probabilities and a logistic regression formula – programmed into a pocket calculator – with which the physician can compute the probability that a patient has acute ischemic heart disease and therefore requires the coronary care unit. Physicians, however, typically do not understand logistic regression and are not happy using this and similar complicated systems (Pearson et al., 1994). The dilemma the hospital now faced was as follows: Should patients in life-and-death situations be classified by intuitions that are natural but sub-optimal, or by complex calculations that are alien but possibly more accurate? This dilemma arises in many contexts, from financial advising to personnel recruiting: Should we rely on experts' intuition or on a fancy statistical model?

Once again there is another alternative: smart heuristics. Green and Mehr (1997) designed a fast and frugal tree by using three building blocks for ordered search, a fast

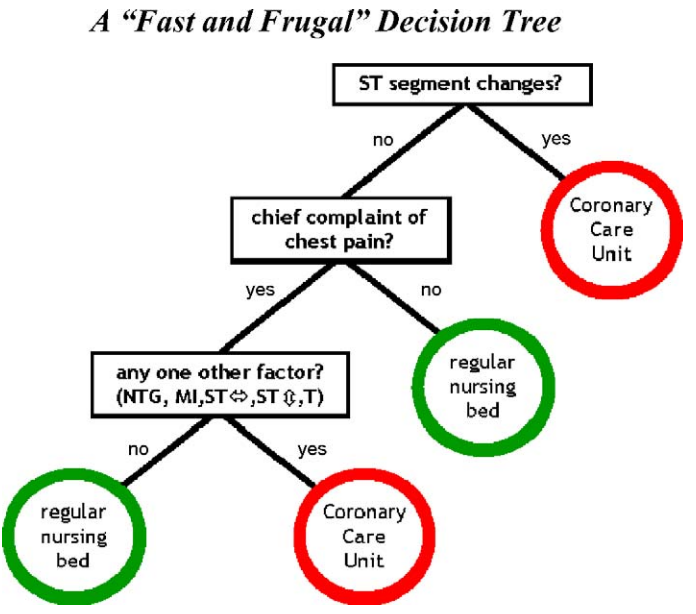

Figure 2. A fast and frugal tree for deciding whether a patient should be placed in a coronary care unit or a regular nursing bed, depending on as few cues as possible. At each stage in the tree, a single cue (in the rectangular boxes) determines whether an immediate decision (in the circles) can be made or if more information must be checked.

stopping rule, and one-reason decision making. The resulting heuristic is shown in Figure 2. It ignores all 50 probabilities and asks only a few yes-or-no questions. If a patient has a certain anomaly in his electrocardiogram (the so-called ST segment change), he is immediately admitted to the coronary care unit. No other information is searched for. If that is not the case, a second variable is considered: whether the patient's primary complaint is chest pain. If not, he is immediately classified as low risk and assigned to a regular nursing bed. No further information is considered. If the answer is yes, then a third and final question comprising additional cues is asked to classify the patient.

This decision tree employs fast and frugal rules of search, stopping, and decision. First, the predictors are looked up in an order determined simply by their sensitivity and specificity. No attempt is made to compute an "optimal" order, such as using conditional probabilities or beta weights, and dependencies between cues are disregarded. Second, search can stop after each predictor; the rest will be ignored. Third, only one predictor determines each decision. The heuristic does not combine – weight and add – the predictors; for instance, a change in the ST segment cannot be compensated for by any of the other predictors. This decision rule is thus an instance of one-reason decision making. The entire heart disease tree is a realization of a fast and frugal tree, which is defined as a decision tree with a small number of binary predictors that allows for a final decision at each branch of the tree.

The simple tree can be evaluated by multiple performance criteria. A good decision strategy (a) has a high sensitivity, that is, sends most of the patients who will actually have a serious heart problem into the coronary care unit; (b) has a low false alarm rate, that is, sends few patients into the care unit who will not need it; (c) is fast in situations where slow decision making can cost a life; (d) is frugal, that is, able to make good decisions with only limited information; and (e) is transparent, so that it will be accepted by physicians and actually be used. The fast and frugal tree is, by design, superior in speed, frugality, and transparency to the decision-support system with its logistic regression and 50 probabilities. But how accurate is it? The counterintuitive result is that the fast and frugal tree was more accurate in classifying actual heart attack patients than both the physicians' intuition and the Heart Disease Predictive Instrument. It correctly assigned the largest proportion of patients who subsequently had a myocardial infarction into the coronary care unit. At the same time, it had a comparatively low false alarm rate. Note that the expert system had more information than the smart heuristic, and could make use of sophisticated statistical calculations. Nevertheless, in this complex situation, less is more. Simplicity can pay off.

5. Ecological Rationality

How can heuristics that ignore information nevertheless be accurate? First, heuristics can exploit environmental structure. As shown by the gaze heuristic, their rationality is ecological. In the following chapters, we give examples of two classes of heuristics that exploit different ways that information can be structured in an environment to help them make accurate decisions. Heuristics that employ ignorance-based decision making (Chapters 106 and 107) exploit environments in which missing data (in particular, lack of name recognition) occur not randomly but systematically. The use of these heuristics is ecologically rational when lack of recognition (e.g., of brand names, stocks, sports teams) is correlated with the criterion (e.g., quality, market performance, winning a game). Heuristics that employ one-reason decision making (Chapter 108) can exploit environments in which the importance (e.g., beta weights in regression) of the cues available are exponentially decreasing, that is, non-compensatory. If this is the case, one can prove that a simple heuristic called Take The Best can perform as well as any "optimal" linear combination of binary cues.

The second reason why simplicity can be smart is its robustness, that is, the ability to generalize well to new environments – specifically to those whose structure is not known in advance. The important distinction here is between data fitting and prediction. In the first case, one fits a model to the empirical data, that is, the training set is the same as the test set. In prediction, the model is based on a training set, but tested on new data. A good fit may be deceptive because of overfitting. In general, a model A overfits the training data if there exists an alternative model B, such that A has higher or equal accuracy than B in the training set, but lower accuracy in the test set. Models with many free parameters, such as multiple regression or Bayesian methods, tend to overfit in environments where information is noisy or fluctuating, particularly when forced to make predictions from small samples.

As an illustration, consider again the coronary care unit allocation problem. The decision-support system using logistic regression was validated on several thousand patients in a large hospital. But it was subsequently applied in different hospitals to new groups of patients who deviated in unknown ways from the original sample. As a result, the model that was best in the original population was no longer guaranteed to be the best in those new situations. There are statistical techniques that expend considerable computational power and time to try to determine the point at which a model maximizes its predictive accuracy without overfitting. Fast and frugal heuristics sidestep this expenditure – their simplicity alone helps them avoid overfitting and perform robustly. For instance, in Chapter 108 we will see that the Take The Best heuristic generalized robustly across 20 predictive tasks, outperforming multiple regression, which overfitted the data.

To summarize, the reasonableness of fast and frugal heuristics derives from their ecological rationality, not from following the classical definitions of rationality in terms of coherence or internal consistency of choices. Indeed, some of the heuristics can produce intransitive inferences in direct violation of standard rationality norms, but they can still be quite accurate (Gigerenzer, Czerlinski, and Martignon, 1999). A heuristic is ecologically rational to the degree it is adapted to the structure of information in an environment, whether the environment is physical or social.

6. What is to Come

The heuristics presented in the coming chapters of this Handbook are examples drawn from a larger group of fast and frugal mental mechanisms that various researchers have begun to explore, both analytically and experimentally. All of these heuristics are process models, allowing for stronger experimental tests than the outcome models that are more typically studied in economics. A variety of experimental set-ups are being developed to test the search, stopping, and decision rules of heuristics, including information-board displays (e.g., Mouselab; Payne, Bettman, and Johnson, 1993) and Internet web sites designed for electronic commerce (Jedetski, Adelman, and Yeo, 2002). The empirical studies typically reveal that there are multiple heuristics (often up to four) that people use in any given task, with some being much more commonly used than others. This variation in the tools drawn from the adaptive toolbox is most pronounced when people are unfamiliar with a task or in environments with flat maxima (where several strategies perform equally well). Given this observed variation, it is necessary to test multiple models of heuristics against one another in a competitive way (rather than performing null hypothesis testing), and to analyze experimental data on an individual rather than aggregate level to discover which heuristic each person is using.

In the next chapter, we introduce the simplest heuristic in the adaptive toolbox, the recognition heuristic, which is an instance of ignorance-based decision making that gives rise to the counterintuitive less-is-more effect. In Chapter 107 we apply the recognition heuristic to a chaotic, unpredictable environment – the stock market – and show how it can guide profitable investment decisions. Chapter 108 introduces one-reason decision heuristics and describes how they fare in comparison with other more traditional inference mechanisms across a variety of environments. In Chapter 109 we discuss the importance of considering environment structure for making cognitive illusions appear or disappear. Finally, in Chapter 110 we present a variety of social heuristics that are designed to promote rapid and appropriate decisions in encounters with other social agents.

These chapters are meant to encourage a vision of economic agents that is anchored in the psychological possibilities of actual humans rather than in the fictional construction of economic optimizers. The widespread opposition between the psychological and the rational is, in our view, an unnecessary fiction. The real human mind can be ecologically rational – its heuristics can be fast, frugal, and accurate at the same time.

References

Berretty, P.M., Todd, P.M., Martignon, L. (1999). "Categorization by elimination: Using few cues to choose". In: Gigerenzer, G., Todd, P.M., the ABC Research Group (Eds.), Simple Heuristics that Make Us Smart. Oxford University Press, New York, pp. 235–254.

Camerer, C. (1998). "Bounded rationality in individual decision making". Experimental Economics 1, 163–183.

Elman, J. (1993). "Learning and development in neural networks: The importance of starting small". Cognition 48, 71–99.

Friedman, M. (1953). "Essays in Positive Economics". University of Chicago Press, Chicago.

Gigerenzer, G. (1996). "On narrow norms and vague heuristics: A reply to Kahneman and Tversky". Psychological Review 103, 592–596.

Gigerenzer, G., Selten, R. (Eds.) (2001a). Bounded Rationality: The Adaptive Toolbox. MIT Press, Cambridge, MA.

Gigerenzer, G., Selten, R. (2001b). "Rethinking rationality". In: Gigerenzer, G., Selten, R. (Eds.), Bounded Rationality: The Adaptive Toolbox. MIT Press, Cambridge, MA, pp. 1–12.

Gigerenzer, G., Czerlinski, J., Martignon, L. (1999). "How good are fast and frugal heuristics?" In: Shanteau, J., Mellers, B., Schum, D. (Eds.), Decision Research from Bayesian Approaches to Normative Systems: Reflections on the Contributions of Ward Edwards. Kluwer, Norwell, MA, pp. 81–103.

Gigerenzer, G., Todd, P.M., the ABC Research Group (1999). "Simple Heuristics that Make Us Smart". Oxford University Press, New York.

Green, L., Mehr, D.R. (1997). "What alters physicians' decisions to admit to the coronary care unit?" The Journal of Family Practice 45, 219–226.

Hertwig, R., Todd, P.M. (2003). "More is not always better: The benefits of cognitive limits". In: Macchi, L., Hardman, D. (Eds.), Thinking: Psychological Perspectives on Reasoning, Judgement and Decision Making. Wiley, Chichester, UK.

Holton, G. (1988). "Thematic Origins of Scientific Thought", second ed. Harvard University Press, Cambridge, MA.

Jedetski, J., Adelman, L., Yeo, C. (2002). "How web site decision technology affects consumers". IEEE Internet Computing 6, 72–79.

Kahneman, D., Tversky, A. (1979). "Prospect theory: An analysis of decision under risk". Econometrica 47, 263–291.

Kahneman, D., Tversky, A. (1996). "On the reality of cognitive illusions". Psychological Review 103, 582–591.

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.

Payne, J.W., Bettman, J.R., Johnson, E.J. (1993). "The Adaptive Decision Maker". Cambridge University Press, New York.

Pearson, S.D., Goldman, L., Garcia, T.B., Cook, E.F., Lee, T.H. (1994). "Physician response to a prediction rule for the triage of emergency department patients with chest pain". Journal of General Internal Medicine 9, 241–247.

Polya, G. (1954). "Mathematics and Plausible Reasoning: Induction and Analogy in Mathematics", vol. 1. Princeton University Press, Princeton.

Pozen, M.W., D'Agostino, R.B., Selker, H.P., Sytkowski, P.A., Hood, W.B. (1984). "A predictive instrument to improve coronary-care-unit admission practices in acute ischemic heart disease". The New England Journal of Medicine 310, 1273–1278.

Rabin, M. (1998). "Psychology and economics". Journal of Economic Literature 36, 11–46.

Sargent, T.J. (1993). "Bounded Rationality in Macroeconomics". Oxford University Press, Oxford, UK.

Selten, R. (2001). "What is bounded rationality?" In: Gigerenzer, G., Selten, R. (Eds.), Bounded Rationality: The Adaptive Toolbox. MIT Press, Cambridge, MA, pp. 13–37.

Simon, H.A. (1955). "A behavioral model of rational choice". Quarterly Journal of Economics 69, 99–118.

Simon, H.A. (1990). "Invariants of human behavior". Annual Review of Psychology 41, 1–19.

Todd, P.M., Gigerenzer, G. (2000). "Précis of 'Simple heuristics that make us smart'". Behavioral and Brain Sciences 23, 727–741.

Tversky, A. (1972). "Elimination by aspects: A theory of choice". Psychological Review 79, 281–299.