# CHAPTER 1
# Introduction

## 1.1 Goal and Content of the Dissertation

The focus of this dissertation is a recent approach to the study of human decision making that questions how far classical notions of rationality should enter a study of the human mind. Rather than viewing the human reasoning faculty as consisting of domain-general mechanisms that—ideally—should obey normative yardsticks set by logic, probability and rational choice theory, this alternative approach holds that the human mind is equipped with a toolbox of specific mechanisms that work well in particular environments (but fail in others). In their *fast and frugal heuristics* program, Gigerenzer, Todd, and the ABC Research Group (1999) proposed a set of decision strategies that can be simple and accurate at the same time by exploiting (a) fundamental cognitive abilities (e.g., the capacity for recognition) and (b) systematic patterns in the environment. Due to this fit with the environment, these cognitive tools exemplify what Gigerenzer and colleagues call *ecological rationality*, which is defined as "adaptive behavior resulting from the fit between the mind's mechanisms and the structure of the environment in which it operates" (Todd & Gigerenzer, 2000, p. 728). The term 'environment' is meant here to refer to the statistical properties of a decision task, that is, the distribution of the target variable, cue validities, and cue intercorrelations. The notion of ecological rationality binds together the empirical work reported here and serves as the theoretical framework for this dissertation. By examining previously studied mechanisms and applying the framework to a new domain I both critically test central claims of the fast and frugal heuristics program and explore its merits for understanding human decision making more broadly.

In general, I am interested in how we deal with information in memory that represents samples of the environment and as such reflects aspects of the environment. This reflection allows us to use such samples to infer latent properties of the environment. My examination of ecological rationality concerns two types of information in memory. First, I consider recognition memory—the ability to discriminate between previously encountered and novel objects. Whether we have encountered an object before or not is often systematic rather than random, and this nonrandomness can be exploited for making inferences about the objects. Second, I consider how information about people in our social networks is used for inferences. As the people we know represent a sample of the general population, our knowledge of these people can be useful to make inferences about the population.

Judgments based on these two types of knowledge are described in two simple inference principles, which structure the dissertation in its two main parts. In the first part, the focus is on the *recognition principle*, and in particular, I examine the empirical use of this heuristic inference principle. By connecting it to the memory literature, I attempt to reveal more about the role of recognition knowledge in comparison to additional knowledge when people make inferences. Moreover, I examine the conditions under which the counterintuitive *less-is-more effect*, occurs as a result of using the recognition principle.

In the second part, I propose, elaborate and test the *recall principle*, which refers to the use of recalled instances from a person's social environment to make inferences about the population. This notion was initially described as one of the possible processes underlying the availability heuristic (Tversky & Kahneman, 1973), and here it is elaborated as a detailed process model. Moreover, the recall principle is contrasted with an alternative approach, according to which inferences are made by use of general characteristics about an object (rather than idiosyncratic samples).

Both inference principles thus describe the use of samples, as both recognition and recall reflect a sample of objects (or people) from the environment. This dissertation examines how well these samples reflect the world and in what way they can and are being used in the inferences we make about it.

In the next sections I elaborate on the concept of ecological rationality and illustrate some of its key arguments by contrasting a selection of ecological approaches in psychology. But note that the relationship between cognition and the environment has also been looked at from, for instance, philosophical (e.g., Clark, 1997) and engineering (e.g., Brooks, 1999; Clancey, 1997) perspectives.

## 1.2 Lenses, Mirrors and Scissors: What is Ecological Rationality?

What constitutes a good judgment? To be rational, normative frameworks of classical rationality such as logic, probability and rational choice theory require reasoning to comply with *coherence* criteria (for an overview see Baron, 2004; Over, 2004). For instance, preferences must be transitive, beliefs consistent, and probability judgments additive. Ecological rationality, in contrast, is concerned with the actual success of decisions in the real-world. This success depends on the structure and requirements of a given environment (including situational factors such as time pressure), and can, for instance, be measured in terms of accuracy, frugality and speed. These are called *correspondence* criteria (Hammond, 1996), and complying to correspondence criteria will not infrequently be at odds with coherence criteria—and vice versa. For instance, a sports bettor might expect team A to win against team B and team B against team C, but think that team C will win against team A and, though being intransitive, be perfectly accurate.

It is worth stressing at the outset that ecological rationality as described by Gigerenzer et al. (1999) has (at least) two dimensions and can refer to both a decision mechanism and a decision maker. A decision mechanism is called ecologically rational if it exploits particular information patterns in the environment. But using such a mechanism does not automatically render a decision maker ecologically rational. Rather, the decision maker must apply the mechanism in the environments to which it is matched. Concerning this latter dimension of ecological rationality, one might debate whether the term 'adaptivity' is not equivalent or even more appropriate (cf. Payne, Bettman, & Johnson, 1993). Bröder (2005; see also Oberauer, 1998) proposes to keep the terms of rationality and adaptivity separate as they touch different aspects of behavior: In his view, adaptivity refers to the functionality of behavior (e.g., whether it leads to a desired goal), and he argues that rationality demands not only adaptive behavior, but also that it is rule-based, explicit and consistently applied.[1] To avoid confusions, I will stick to Gigerenzer et al.'s use of the term ecological rationality in this dissertation. Both the ecological rationality of inference mechanisms and of the decision maker will be addressed, but the overall emphasis is clearly on former.

The idea that cognitive mechanisms are matched to the informational structure in the environment has been taken up more systematically in decision making only recently. It is, however, by no means new in psychology. For instance, it has been an influential approach in early theories of perception. Therefore, in the next sections I review some prominent examples of the notion that the mind is shaped by its interplay with the environment.

This overview has three aims. First, it is intended to delineate the roots of the fast and frugal heuristics program. Second, it will demonstrate the range of solutions to how the study of the mind can be combined with the study of the environment—and where the different solutions converge and diverge. Lastly, the overview will help us map the position of the fast and frugal program in a broader context. Due to these specific aims, I limit the discussion of the different approaches to those aspects that are relevant in the context of the fast and frugal heuristics, and I hope not to distort the essence of their unique contribution.

I will start with two examples from perception, as it has been in this field of psychological enquiry that the joint study of cognitive mechanism and the environment has been recognized most readily. This will be followed by more general frameworks for the study of cognition, and I will then turn to how the idea has been developed in the field of judgment and decision making—which is the focus of this dissertation. Finally, in light of these precursory approaches, I describe the fast and frugal heuristics program in more detail.

## 1.2.1 Ecological Pioneering in Perception: Brunswik and Gibson

### 1.2.1.1 Egon Brunswik

Egon Brunswik's *probabilistic functionalism* constitutes one of the central theoretical columns on which the fast and frugal heuristics program rests (Gigerenzer, Hoffrage, & Kleinbölting, 1991). Originally, Brunswik developed his framework in the context of perception but he also recognized its relevance in understanding motivation and action (later, Hammond applied it to judgment; see section 1.2.3). The term 'functionalism' indicates that the framework is concerned with how the interplay between means—for instance, actions, or

[1] It should be noted, however, that according to this definition of rationality some of the mechanisms studied by the fast and frugal heuristics program actually exemplify (some sort of) rationality (rather then simply adaptivity).

cues—is organized with reference to some end, or objective, and thus is *functional* for this objective. The key aspect, however, lies in the assumed *probabilistic* nature of the relationship between means and ends. For instance, with perception, the patterns of stimulation on the sense receptors caused and the properties of the world are only probabilistically related.[2] As an organism is often "[f]orced to react quickly and within reasonable limits of time", it is necessary to "respond before direct contact with the relevant remote conditions in the environment" can be made (Brunswik, 1952, p. 22). As a consequence, an organism typically does not have direct access to the world, leaving it with the uncertainty between the *proximal cues* available and the *distal variables* of which the proximal cues are indicators. The central question that Brunswik concerned was how organisms deal with the constraints imposed by this uncertainty, in particular, how they manage to achieve perceptual stability—the central 'achievement' of perception.

One requirement for achievement is that the available information is used effectively, that is, in accordance with its value to achieve the goal. To describe this value, Brunswik introduced the concept of *ecological validity*, which he defined as the correlation coefficient between a proximal cue and the distal criterion. *Functional validity*, conversely, describes the extent to which a cue is actually used and is defined as the correlation between proximal cue and judgment. The intercorrelations between the criterion, cues and judgment are summarized in Brunswik's famous *lens model*. Ideally, ecological and functional validity should be matched, in other words, the way in which an organism uses a cue should be adapted to the environment. Brunswik assumed that perceptual stability in a dynamic environment (in which both the ecological validities and the availability of the cues change) is achieved through a process he called *vicarious functioning*, by which he meant that cues are mutually substituted for each other. As cues are intercorrelated (due to their common dependency on the criterion which often causes them) they are characterized by *equipotentiality*. To wit, Brunswik conceived an organism as an intuitive statistician, who aggregates, substitutes and combines cues, allowing it to reach stability in the face of the uncertainty and dynamic in the environment.

The implications of Brunswik's approach is that to understand achievement, both the environment (texture in the environment in terms of the ecological validities of the cues) and the mechanisms used by the organism have to be studied.[3]

A study by Brunswik and Kamiya (1953) illustrates the environmental approach and highlights its potential in understanding cognitive mechanisms. Analyzing stills from Alec Guiness' film "Kind Hearts and Coronets", the authors provided an ecological argument for the Gestalt factor of proximity, as there was a correlation between the proximity of two

---

[2] This aspect of Brunswik's theory was hard for him to promote in a time in which the dogma of determinism that the behaviorists followed ruled.

[3] Though Brunswik became interested in actually studying the processes underlying achievement only later in his career.

parallel lines in pictures and the presence of an object in the scene. The cognitive law of proximity thus seems to represent a match to an actual regularity in the world.

One important consequence of Brunswik's probabilistic functionalism for empirical research is his call for *representative design* in psychological experiments (Brunswik, 1955; see Dhami, Hertwig, & Hoffrage, 2004, for a recent overview). As the mind, in particular the process of vicarious functioning, is adapted to environmental texture (i.e., cue validities and cue intercorrelations) it has to be studied in the context people are typically dealing with, as otherwise the cognitive function might not be able to display its full potential. Though it is still discussed how exactly representative design should be realized (or checked whether it was achieved), there have been attempts to implement it in psychology (e.g., Gigerenzer, Hoffrage, & Kleinbölting, 1991; see Dhami, Hertwig, & Hoffrage, 2004, for an overview) and the evidence suggests that it can indeed lead to different conclusion concerning the performance of the human mind.

#### 1.2.1.2 James Gibson

A second influential approach to perception putting a strong emphasis on the interplay between the mind and the environment is James Gibson's theory of *direct perception* (Gibson, 1966; 1979). As Brunswik, Gibson was interested in the study of achievement (i.e., veridical perception), and he assumed this is achieved by the exploitation of systematic structures in the information from the environment hitting the retina. One type of information that Gibson emphasized was light. Light is especially rich as it compresses a large amount of information about the environment. To illustrate, as light travels in straight lines light reflected from objects in the environment reveals a lot about these objects (e.g., their substance, position etc)—if it is properly decoded. Another aspect of the environment that Gibson elaborated was the role of *invariants*, lawful patterns of change which the perceptive system can extract. Examples for such invariants are the systematic decrease or increase in changes of the texture gradient, which allows for an inference about the position and the tilt of a surface in space; or transformational invariants, that is, systematic ways in which the apparent area of an object changes as it moves away from the spectator at constant speed.

Veridical perception, according to Gibson, is achieved by *resonance* with these systematic information patterns. Although Gibson used the term more restrictively in later writings, the functions that the environment offers to an organism, and thus also the value invariants have for the perception of the environment (in particular those characteristics that are significant for the organism), can be described as *affordances*. Like Brunswik, Gibson pointed out that veridical perception and effective behavior require *attunement* to the functional values (i.e., the affordances) of the objects in the environment. Yet unlike Brunswik's emphasis on the probabilistic relationship between the information available to an organism and the environment, Gibson argued that invariants uniquely specify the environment, and thus no probabilistic inference is required. Rather, invariants are simply "picked up" and lead directly and unequivocally to the perception of the properties of the distal world (conditional on attunement being achieved). In other words, Gibson's concept of direct perception is at odds with Brunswik's tenet of perception mediated by proximal cues in the environment.[4] This assumed direct relationship later led Gibson to the extreme conclusion that the main focus should be on the study of the environment.

### 1.2.2 Ecological Approaches to Cognition: Shepard and Anderson

#### 1.2.2.1 Roger Shepard

Like Gibson, Roger Shepard (e.g., 1987, 1994) pursued the idea that cognition is tuned to invariant patterns in the environment. Shepard assumes that if relevant, accessible and permanent, such invariants became internalised over the course of evolution and now find their expression in abstract general principles in cognition. These internalized universal regularities help the mind disambiguate situations that would otherwise be unsolvable. But in contrast to Brunswik and Gibson, Shepard reached considerably beyond perception and searched for such reflections of the environment in cognitive functions more generally, including higher-order ones.

Take categorization, that is, how the human mind generalizes across objects, as an example. Shepard (e.g., 1958) showed that people collapse similar objects into one category if they fall into a "consequential region", following a particular, constant gradient of generalization. This gradient, so the argument goes, constitutes an internalization of an environmental regularity "because the set of objects having a particular consequence will generally be physically similar to each other, the set will constitute a compact region … in the abstract space of the objects" (1987, p. 265). As another example, Shepard argues that the apparent trajectory of an object when seen successively at two positions (i.e., the way representations of objects are rotated mentally) is due to internalized constraints reflecting the possible projections and transformation of an object in three-dimensional Euclidean space (condensed in the principles of kinematic geometry; e.g., Shepard, 1984).

Note that in contrast to Brunswik, who argued that relevant aspects of the world are reconstructed by an inference process, Shepard assumes the structure of the environment to be internalised apriori. Similarly, Shepard (as Gibson) is in opposition to Brunswik's emphasis on uncertainty as an irreducible property of the environment, as the assumed internalization of the environment dissolves possible uncertainties. Importantly, possibly because he also studied higher-order processes, Shepard—more than Brunswik and Gibson—acknowledges that cognitive mechanisms are subject to the constraints of speed and simplicity. For instance,

---

[4] This hold in spite of the fact that Gibson also speaks about the "specification" of the environment by picked up info, and thus seems to acknowledge indirectly some indirectness in the relationship between information and environment. Generally, however, note that what Gibson rejected was less the fundamental distinction between the environment and the information available to an organism about that environment, but rather the empiricist position that rests on the assumption that perception relies on an elaboration of sensations by inferential processes utilizing memory, habit etc. Rather, he argued that there is an immediate relation between the information and the perception of the environment, without having to decompose it to sensory elements (cf. Gordon, 1989).

speaking about apparent motion, he argues that the one representation of a motion will be realized that "is the geometrically simplest and hence, perhaps, the most quickly and easily computed" (p. 585). The issue of natural cognitive constraints will become a central one in some of the approaches discussed below.

Finally, concerning methodological implications of Gibson's thesis of a close mind-environment match he argued that "[t]he internalized constraints that embody our knowledge of the enduring regularities of the world are likely to be most successfully engaged by contexts that most fully resemble the natural conditions under which our perceptual/representational systems evolved" (1987, p. 266), resonating Brunswik's call for representative design.

#### 1.2.2.2 John Anderson

A concrete methodology for developing theories of cognitive processes that are adapted to the environment was proposed by John R. Anderson (1990, 1991). The starting point of his *rational analysis* (see Chater & Oaksford, 1999, for a recent review) is the assumption that human cognition represents an optimal response with regard to the goals of the organisms and the structure of then environment—within the bounds of an organism's natural capacities (putting an upper limit to the complexity of possible solutions). Take, for instance, the structure of memory. As retrieval of memory traces is constrained by the cognitive costs it produces, the retrieval costs have to be balanced against the expected gain from the additional information for a given task weighted by the probability that a memory will be relevant in the current context. Crucially, this probability is mainly a function of the structure of the environment, in particular the rate with which particular traces have been relevant in the past. As Anderson and Schooler (1991) could show both how the probability that an item occurs in the environment (holding context constant) and how the decay of traces in memory increases with increasing retention interval (i.e., the time that the items was used last) can be described by a power function. This parallel suggests that retention in memory is matched to the pattern in which items occur in the environment. Thus, memory appears to represent an optimal response to the structure of the environment conditional on the natural cognitive bounds. If the two functions did not match, memory would not be optimal as the costs would not be optimally related to gains.

Based on the approach of rational analysis, there have been successful attempts to reinterpret human behavior previously viewed as irrational as an optimal response given a particular environmental structure (e.g., Oaksford & Chater, 1994). For instance, biases in covariation assessment (e.g., McKenzie & Mickkelsen, in press) or the "confirmation bias" in the Wason selection task appear to be rational responses under minimal and reasonable assumptions about the environment (e.g., that most events are rare).

In comparison to the approach to ecological rationality taken by the fast and frugal heuristics program (outlined in more detail below), it is important to point out that rational analysis retains the assumption that the goal of an organism is to produce an optimal response. However, it remains unspecified what processes might achieve this optimizing.[5]

After the last two approaches we considered claim to provide approaches to cognition generally, we now turn to approaches that more specifically aim at instantiating an ecological perspective in models of judgment and decision making—the topic of this dissertation.

### 1.2.3 Judging In Tune with the Environment: Social Judgment Theory

Kenneth Hammond (1955), a student of Brunswik, applied Brunswik's probabilistic functionalism to the study of human judgment. Note that although Brunswik initially reserved the 'uncertainty-geared' processes described by probabilistic functionalism to mechanisms in perception, and distinguished them from the 'certainty-geared' processes of reasoning, he later acknowledged the probabilistic nature of many higher order cognitive functions.[6] To emphasize these commonalities, he coined the term 'ratiomorphic' to refer to processes sharing this probabilistic character.

Hammond's initial application was in the context of clinical judgments. That is, he was concerned with modelling judgments concerning diagnosis and prognosis, an area that aroused enormous interest after Meehl's (1954) comparison of clinical and statistical prediction (see Goldstein, 2004, for a recent overview). By applying the Brunswikian perspective to judgment and by elaborating the lens model, Hammond provided a methodology to study human judgment systematically. In particular, the lens model approach allowed for a simultaneous study of the relationship of the variables in the environment (in the sense of the relationship between cues, judgment, and the environment). Reflecting Brunswik's focus on achievement, Hammond's Social Judgment Theory evaluated judgments using correspondence rather than consistency criteria. The tool used for 'policy capturing', the analysis of the weighting schemes followed by a judge, is typically multiple regression. While in the beginning multiple regression represented merely a research tool and paramorphic model of the judgment process (i.e., only modelling the end product of a decision mechanism rather than the actual computational steps), it was later also interpreted as an adequate model of the mind (Brehmer, 1994).

[5] In addition, it can be debated whether the Bayesian framework, on the basis of which an optimal response is defined, is an appropriate normative system (Gigerenzer, 1991). After all, for a given situation different norms can be brought to bear (Gigerenzer, 1996).

[6] Brunswik distinguished uncertainty-geared processes from 'certainty-geared' processes as the former aim for "smallness of error at the expense of the highest frequency of precision", whereas the latter strive for the highest accuracy. As Brunswik pointed out, the two can be distinguished by different error distributions: the former will produce few perfectly accurate estimates and a nearly continuous distribution of errors with small variance, whereas the latter has a potential to be very accurate and consistent, but also to produce systematic and large deviations, apparent in a very irregular discrete distribution of errors with a large variance. To wit, the uncertainty-processes forsake highest precision for an overall smallness of error, foreshadowing Simon's notion of 'satisficing'.

### 1.2.4 Scissors and Bounds (and Biases): Heuristic Decision Making

Due to its roots in perception, where information acquisition typically does not produce much cognitive costs and where many variables can be processed simultaneously, the Brunswikian approach tends to put little weight on the cognitive bounds that an organism faces. To be sure, probabilistic functionalism acknowledges bounds, but these bounds are due to the limited availability and imperfect ecological validity of the proximal cues, limiting the accuracy achievable by an organism. Yet, in particular for higher-order cognitive functions such as judgment and decision making, can it be expected that always all information is considered? Can real people be expected to perform the weighting and substituting underlying vicarious functioning? Moreover, how is information searched, when, in contrast to perception, information has to be searched actively and processed sequentially? And when is search stopped?

One approach focussing on the natural limits of real people under realistic situations is Herbert Simon's concept of *bounded rationality*. This approach stresses that traditional models of rationality seem like an ill-fit idealisation. Due to internal (e.g., memory) and external (e.g., incomplete information, time pressure) constraints, people often have to use heuristic strategies when making decisions. But this is not the end of the story. Simon added to this insight a by now familiar argument: that to understand human rational behavior it is important to study the environment (Anderson, for instance, explicitly refers to Simon). However, Simon's unique contribution was to point out that the match between the mind and the environment might have the function of simplifying the cognitive processes involved. To describe this interplay, Simon (1990, p. 7) proposed the metaphor of "a scissors whose two blades are the structure of the task environments and the computational capabilities of the actor". To understand rational behavior it is important to study the structure of the environment, as the environment might indicate in what respect a mechanism was simplified—or was not pressured to evolve into a more complex one. Ideally, the simplification that the environment invites coincides with a cognitive bound. Due to such a fit, what looks simplistic in isolation might turn out to be an intelligent adaptation.

Another important contribution was to point out that for most tasks it is not necessary—and often simply impossible due to computational intractability—to find the globally optimal option (as assumed in many standard economic theories). Instead, a good enough solution is typically sufficient. Simon coined the term *satisficing* for aiming at a just good enough solution.

Simon's concept of bounded rationality was later invoked as one of the starting points for the *heuristics and biases program* (e.g., Kahneman, Slovic, & Tversky, 1982; Kahneman, 2003). This program introduced the idea that when making judgments, rather than trying to emulate normative procedures based on logic, probability or rational choice theory, people use simplifying heuristics that operate on very little and easily accessible information without integrating all relevant variables. Although these substitute mechanisms often produce useful solutions, in some environments they can produce systematic errors, and the program typically studied heuristics in these environments. As a consequence, the main message of the heuristics and biases program was often reduced to showing that the human mind is irrational (cf. Krueger & Funder, 2004). It has to be noted, however, that this take on the bounds of human reasoning distorts Simon's original concepts, as his goal of introducing the concept of bounded rationality was to criticize the norms typically applied to evaluate the behavior, rather than stressing people's irrationality. Kahneman and Tversky, however, stick to the norms of classical rationality (i.e., probability theory, logic, rational choice theory), the appropriateness of which remains highly controversial (e.g., Gigerenzer, 1996; Kahneman & Tversky, 1996; for an overview, see McKenzie, 2005).

One unfortunate aspect—from an ecological perspective—of the heuristics and biases program is that typically no attempt is undertaken to analyse the conditions under which a heuristic works well.[7] Therefore, the heuristics and biases program does not pay attention to the way in which the environment might shape human behavior, but instead focuses on situations in which there is a mismatch between a heuristic and the environment. To be sure, the study of errors might sometimes be a useful research strategy (just think of the insights gained by the study of optical illusions). However, in light of the ecological approaches in perception described above, such an approach risks to underestimate the functional value of the cognitive mechanisms. Moreover, by generally neglecting such environmental issues and not specifying the antecedent condition that give rise to the heuristics proposed, ironically the heuristics and biases program seems to have in a sense inherited the domain-general character from the statistical all-purpose models they criticized (Bayesianism, multiple regression; e.g., Edwards, 1968; Peterson & Beach, 1967).

### 1.2.5 Summary

The preceding overview has taken us from ecological approaches in perception to reasoning to decision making. It should be emphasized that one can debate about whether the same general arguments for an ecological perspective apply across these different cognitive functions. For instance, invariant structures might be faced less often in decision making than in perception, and reasoning and decision making may be more 'certainty-geared'—aiming at the highest possible accuracy—than perception. One possible consequence of such differences might be that for higher order cognitive processes there are multiple ways to solve a cognitive task, which is acknowledged in the many 'dual systems' approaches to reasoning (e.g., Evans & Over, 1996; Sloman, 1996; Stanovich & West, 2000; Petty & Cacioppo, 1986; see also Hammond, 1996). Nevertheless, characteristics such as uncertainty in the environment and the involvement of mediating information between mind and environment are often present across all cognitive functions, and thus should have had an impact on their evolution.[8]

---

[7] This problem is amplified by the fact that the heuristics were not clearly formulated, and no antecedent conditions for their use were specified (cf. Gigerenzer, 1996).

[8] Note that this is even true for Gibson, the proponent of *direct* perception (see Footnote 4).

In general, during the overview we encountered different solutions to describe the mind-environment interaction, differing, for instance, in terms of the acknowledgement of the limits of the cognitive system, the role of optimality, and the models used to describe the underlying processes. In addition, the different approaches reviewed use different metaphors for describing the interplay between mind and environment. For Brunswik, it was a lens that best captured this relationship; Shepard chose a mirror, and Simon a pair of scissors. Which one is most useful? In contrast to mirrors, the scissors metaphor implies that there is no exact correspondence between mind and environment, but only aspects of environment relevant for a particular task are attended to—the rest is ignored. The lens metaphor, in contrast, implies that a reflection can only be indirect and mediated by intervening variables, and that the environmental structure is actively reconstructed, rather than internally ingrained. As pointed out by Todd and Gigerenzer (2001), an answer to the question of which metaphor is most apt should depend on the cognitive task considered. For higher cognitive functions such as judgment and decision making, for which simplifying heuristic mechanisms are arguably the rule rather than the exception, one should probably expect less mirror-like reflections, but rather an exploitation of external structures necessary (see also Barlow, 2001). Why? Because the minimization of the amount of considered information that might give rise to such heuristics could be implemented by focussing only on the useful aspects of the environment, rather than using a high fidelity replica of it. (This notion will be elaborated more in the next section.)

Having acquired an overview of various approaches to the notion that the mind is designed as a response to the structure of the environment, we are now in the position to map the specific notion of ecological rationality formulated in the fast and frugal heuristics program more broadly and address the program in more detail.

## 1.3 Fast and Frugal Heuristics: Models of Ecological and Bounded Rationality

The key idea of the fast and frugal heuristics program can be stated succinctly as follows: simplicity combined with a match to the environment can yield astonishingly accurate judgments. Whereas traditional models of judgment and decision making clearly acknowledge that some ‘fit to the environment’ (in the sense that relevant information must be considered and weighted appropriately) is a precondition for accuracy, simplicity has often been equated with the opposite a good judgment (but see Dawes, 1979; Dawes & Corrigan, 1974). Apart from being due to the undue emphasis on errors by the heuristics and biases program, simple mechanism probably came into disrepute by appearing too inflexible to generally yield good performance. So how could simplicity be rehabilitated?

First, note that the problem of inflexibility arises only if the aim of domain-general mechanism is retained. Domain-specific mechanisms, in contrast, can afford to be simple if they succeed in working well in particular, but ideally common environmental structures.[9] The second argument is based on the fact that most environments in which inference mechanisms have to operate are noisy. That is, only some of the information in the environment is relevant for the future performance of the mechanism. The challenge is to ignore the noise and find the truly predictive signal. And it is in solving this task that Davidian simplicity can bring Goliathian complexity to its knees. By attending to all available information, complex models such as multiple regression tend to overfit. Overfitting means that part of the adjustments is made to noise, which does not prove to generalize to new situations and then decreases performance (cf. Geman, Bienenstock, & Doursat, 1992; Massaro, 1988). By contrast, simple mechanisms avoid overfitting by attending to only the most relevant information and then often end up being more robust than complex models (e.g., Dawes, 1979; Czerlinski, Gigerenzer, & Goldstein, 1999).

Gigerenzer et al. (1999) emphasize the goal of developing simple models as they strive to go beyond paramorphic 'as-if' models of the judgment process. Instead, they aim at describing the computational steps involved (at least as regards the algorithmic level of description) that can actually be carried out by people with limited time, memory and computational capacities. In particular, fast and frugal heuristics specify how information is searched, when search is stopped and how a decision is made based on the acquired information. These steps are viewed as the building blocks of adaptive behavior that can give rise to ecological rationality. With their notion of ecological rationality Gigerenzer and colleagues intend to "bring environmental structure back into bounded rationality" (Todd & Gigerenzer, 2001, p. 730). In other words, in contrast to most of the approaches discussed above, cognitive bounds serve as the starting point rather the influence of the environment (if cognitive bounds were an issue at all; e.g., Gibson or Brunswik).

To illustrate the principles of simplicity and match with the environment, consider the *Take The Best* heuristic (*TTB*; Gigerenzer & Goldstein, 1996). TTB is a fast and frugal heuristic for inferring which of a set of objects, has a higher value on the criterion. Assume, for simplicity, the case of a choice between two objects based on binary information, "cues", with positive values indicating that an object has a high value on the criterion. TTB has a search rule that specifies a cue-wise search according to the cues' validities. A stopping rule specifies that search is stopped as soon as a cue is encountered on which the objects have different values. Only this cue is considered for a decision, and TTB's decision rule specifies that the object with the positive value is inferred to have a higher value on the criterion. Note that TTB is *noncompensatory* as the values on cues not looked up cannot reverse the choice determined by the first discriminating cue. TTB is simple as it only requires memorizing the values of one cue at a time and that it requires no integration of values.

[9] Although the terms 'environment' and 'domain' are usually used to distinguish, respectively, between the statistical structure (e.g., high vs. low cue intercorrelation) and the content (e.g., social exchange vs. foraging) of a decision task, Gigerenzer et al. (1999) often use these terms interchangeably to refer to the statistical structure. Accordingly, I will use 'domain-specificity' to refer to the dependence of a mechanisms on a particular statistical structure in the environment, and not to the content of a task.

But TTB is not only simple but also domain-specific. Its domain specificity becomes apparent when we consider its accuracy in various environments. As Martignon and Hoffrage (2002) have shown, TTB performs as good as compensatory mechanisms (e.g., multiple regression) in environments in which the importance of the cues is distributed such that a more important cue is at least slightly more important than the sum of all less important cues. Moreover, it performs comparatively well in environments in which information is scarce. But domain specificity can be obtained in many ways. Consider other fast and frugal heuristics. The QuickEst heuristic (Hertwig, Hoffrage, & Martignon, 1999) is domain-specific as it excels in (the ubiquitous) j-shaped distributions of criterion values; and the recognition heuristic works well in environments in which there is a strong correlation between recognition and the criterion. All these examples exemplify ecological rationality. Importantly, the environmental characteristics (or some approximation of them) to which the heuristics are matched seem to occur frequently in the real world.

One still open issue concerning fast and frugal heuristics is whether they should be seen as ‘betting’ on particular environmental structures (which might pay out if highly likely that environment—or some approximation of it—is present) and thus should be applied generally or whether they are selected as a consequence of learning (Rieskamp & Otto, 2003). If they are selected from the adaptive toolbox by learning, what is the default strategy before any learning could occur?

How do fast and frugal heuristics relate to the other ecological approaches? In general, it is Simon’s vision of a close coupling of the natural bounds of human cognition and the structure of the environment that Gigerenzer et al. (1999) follow in their endeavour to model human decision making. Moreover, Gigerenzer et al. refer to Simon’s notion of satisficing, as—in contrast to Anderson’s rational analysis—fast and frugal heuristics do not aspire to optimize their performance but accept regular, but small error rates. Moreover, TTB can be seen as taking up Brunswik’s idea of vicarious functioning and equipotentiality of cues in the sense that many different cues could serve as the final basis for an inference. TTB simplifies and provides a process model of vicarious functioning by introducing sequential search and stopping rules and one-reason decision making (cf. Gigerenzer & Kurz, 2001).

## 1.4 Two Principles of Ecological Rationality: Recognition and Recall

In the next sections I describe two principles of ecological rationality that are investigated in this dissertation: (a) the recognition principle and (b) the recall principle. These two inference principles can be seen as complementary, as recall builds upon recognition when the latter does not discriminate (i.e., allow to make an unambiguous prediction). In this dissertation, however, I will address them separately. Both inference principles exemplify ecological rationality as they are based on proximal information that reflects latent properties of the distal environment. It is argued that although being based on easily accessible information and representing highly simple mechanisms, courtesy of this reflection the principles can give rise to accurate assessments of the world.

### 1.4.1 Recognition: Exploiting Partial and Systematic Ignorance

Imagine the following situation: you have to infer which of two diseases is more common in Germany. Unfortunately, you know relatively little about medical issues, and you have heard of only one of the diseases—the other disease is unknown to you. How could you go about making this inference? According to the *recognition principle* (Gigerenzer & Goldstein, 1996) you would make a vice out of your partial ignorance and simply pick the recognized disease. Goldstein and Gigerenzer (1999, 2002) instantiated this principle in the recognition heuristic. The recognition heuristic thus exploits the highly evolved capacity in humans (and other organisms) to discriminate between previously encountered and novel objects. Moreover, the recognition heuristic 'inherits' the robustness that recognition has been shown to exhibit in the face of cognitive damage and mental decline (e.g., Warrington & McCarthy, 1988). (The relationship between the recognition heuristic and the recognition memory literature is discussed in Chapter 2.1.) Most importantly, the recognition principle exploits a systematic pattern in the environment, namely that recognition (or a lack thereof) is correlated with various properties of the world.

Note that Gigerenzer and Goldstein conceptualise recognition as binary: either one recognizes an object or one does not, and it is this assessment on which the recognition heuristic operates. Apart from being unambiguously defined, it is this binary nature of recognition that distinguishes the recognition principle from Tversky and Kahneman's (1973) concept of availability: although recognizing an object can also be associated with an increased availability (in the sense of more fluent processing), availability is continuous. Moreover, availability was defined as involving the imagined or actual retrieval of instances (see, e.g., Tversky & Kahneman, 1973, p. 208), whereas recognition does not require that an instance can be retrieved.

By building upon a highly complex cognitive capacity, the recognition heuristic itself can afford to be simple. A crucial claim by Goldstein and Gigerenzer is that recognition is used in a *noncompensatory* way. This means that recognition—if it discriminates—is the only probabilistic cue considered for the inference and all other possible evidence is ignored, regardless of whether this evidence would support or contradict recognition. Although the term 'noncompensatory' is usually used in the judgment and decision making literature to refer to the use of probabilistic cues (or attributes in the context of preferential choice; e.g., Keeney & Raiffa, 1993; Tversky, 1972), Goldstein and Gigerenzer (2002) appear to refer to information in general: "[I]f one object is recognized and the other is not, then the inference is determined; no other information about the recognized object is searched for and, therefore, no other information can reverse the choice determined by recognition" (p. 82).

This claim of a noncompensatory use of recognition has been repeatedly challenged. In Chapter 2.1, I will critically review the accumulated evidence regarding this claim. In addition, although this evidence suggests that recognition does not generally give rise to noncompensatory decisions, I will test an alternative account for why recognition might have a special status in decision making and I discuss possible factors that foster the noncompensatory use of recognition.

In Chapter 2.2 I will address an interesting consequence of using the recognition heuristic. Consider you have to pick in a paired comparison task involving $N$ objects the one with the larger criterion, and $n$ objects are recognized. There will be a total of $N(N\text{-}1)/2$ pairs, encompassing three different types of comparisons. First, there will be $n(N\text{-}n)$ pairs in which only one object is recognized (RU). This represents the situation where the recognition heuristic is applicable. Assuming the recognition heuristic is used for these types of comparisons (i.e., the recognized objects is always chosen to have a larger criterion value—assuming a positive correlation between recognition and the criterion), the accuracy on these pairs represents the *recognition validity*, denoted by $\alpha$. Second, there will be $\frac{(N-n)(N-n-1)}{2}$ pairs in which none of the objects is recognized (UU). In lack of knowledge about these objects, one has to pick one of them randomly, yielding an expected accuracy of 0.5 on these pairs. Finally, there will be $\frac{n(n-1)}{2}$ pairs in which both objects are recognized (RR). These comparisons can be made based on further knowledge one has about the objects. The observed accuracy on these pairs will depend (mainly) on the validity of this knowledge for the given comparison task, and is denoted by $\beta$, the *knowledge validity*[10] (see Chapter 2.2 for further details).

Importantly, the proportions of these different types of comparisons among all comparisons will vary depending on the number of recognized objects, $n$. Figure 1.4.1a shows the different proportions of the three types of comparisons as a function of $n$. The overall accuracy, across all three types of comparisons, is calculated by

$$f(n) = 2\left(\frac{n}{N}\right)\left(\frac{N-n}{N-1}\right)\alpha + \left(\frac{N-n}{N}\right)\left(\frac{N-n-1}{N-1}\right)\frac{1}{2} + \left(\frac{n}{N}\right)\left(\frac{n-1}{N-1}\right)\beta .$$

Note that thus the overall accuracy is a function of $n$. Interestingly, whenever the recognition validity is larger than the knowledge validity (i.e., if $\alpha > \beta$), there will be a point $n_{lim} < N$ in the curve at which the expected overall accuracy at $n_{lim}$ is *larger* than at $n_{lim} + 1$. In other words, in this situation recognizing a smaller number of objects will lead to more accurate inferences, a phenomenon that Goldstein and Gigerenzer (2002) dubbed the *less-is-more effect*. It is shown in Figure 1.4.1b (with $\alpha = .8$ and $\beta = .6$).

In Chapter 2.2, I will trace the emergence of this effect in an empirical study looking at different groups of people that vary in their amount of knowledge (in terms of the number

[10] But note that also other factors influencing the performance, such as time pressure or cognitive load, will have an impact on the value of the knowledge validity (see Chapter 2.1).

of recognized objects). In particular, this study serves as a critical test concerning the conditions under which a less-is more effect can occur.

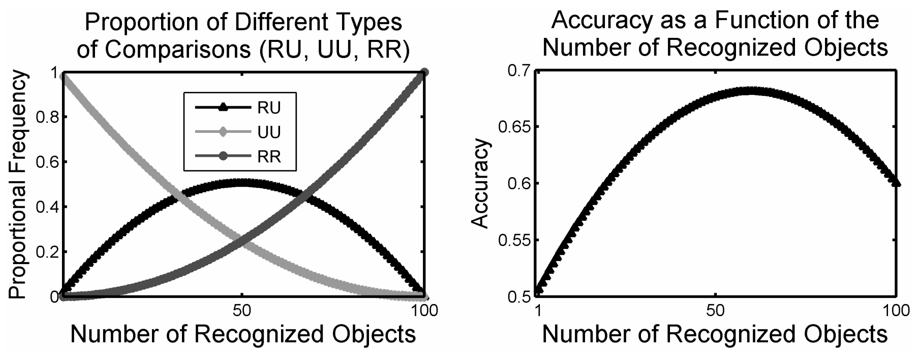


*Figure 1.4.1.* (a) Proportions of the three types of comparisons as a function of the number of recognized objects ($n$) and (b) the less-is-more effect (see text for details).

## 1.4.2 Recall: The Social Environment as the Window to the World

Let us return to the inference task of deciding which of two diseases occurs more frequently in Germany. Whereas the recognition principle is applicable when one disease is recognized but not the other (RU), if $n > 1$ there will also be comparisons in which both diseases are recognized (RR). Here the recognition principle will not be of any help, and one has to recruit further knowledge to make a decision (unless one guesses). One type of knowledge that might then come into play are occurrences of the diseases in one's social environment. In other words, one could base an inference on a sample of the population—one's family, friends, and acquaintances (and oneself).

The use of retrieved instances for an inference about the environment is what I call the *recall principle*, and it applies when different numbers of instances for the events can be retrieved. (Chapter 3.2 considers what knowledge might be used when no or the same number of instances can be recalled and thus the recall principle is not applicable.) The inferences modelled by the recall principle in this dissertation concern the frequency of events in a population. In particular, it concerns events of which people can be manifestations (e.g., names, professions, diseases, leisure activities etc). Similarly as with the recognition principle (where a lack of recognition is taken to be informative), a lack of recall—that is, that one cannot retrieve an instance of one of the events—is informative. In contrast to recognition, however, recall is graded rather than binary, and it can thus make a prediction even when for both events instances can be retrieved (as long as the numbers of instances differ). Thus, if you can recall from your social environment three individuals who had one of the diseases, but only one individual for the other disease, according to the recall principle you simply infer that the disease for which you could recall more individuals in your social network is also more prevalent in the population overall.

Information obtained about concrete individuals from one's social network could receive special prominence for several reasons: First, as no mediating factor can distort it, information directly obtained or observed about the members of one's own social network is highly reliable. Second, the observations of instances of the target criterion are per se a valid indicator of the criterion. Further, the information is easily accessible, as information about social network members represents a constantly recurring and thus well-rehearsed event. Finally, observations of criterial events in one's social environment are naturally sampled, that is, encountered sequentially and represented as natural frequencies, a format that has been shown to foster probabilistic reasoning (Hoffrage, Lindsey, Hertwig, & Gigerenzer, 2000).

As the recognition principle, the recall principle exploits a systematic pattern in the environment. Specifically, as events that are more frequent in the overall population are also more likely to be encountered in one's limited social sphere, frequencies of different events among the people one knows can be used to make an inference about the frequencies of the events in the population.

The notion of frequency judgments based instances retrieved from one's acquaintances has been formulated by Tversky and Kahneman (1973) as one of the mechanisms underlying the availability heuristic. However, whereas this heuristic does not address the question of how instances are sampled, with the recall principle I examine possible ways to describe this sampling process. Moreover, treating the recall principle from the perspective of ecological rationality, I analyze to what extent the principle allows to infer distal properties of the environment. Further differences between the recall principle and the availability heuristic are discussed in Chapter 3.2.

## 1.5 Organization of the Dissertation

The dissertation is organized as follows. Chapter 2 deals with the recognition principle and consists of two parts. Chapter 2.1 is an investigation of the empirical use of the recognition heuristic, and in particular of the question whether recognition has a 'special status' in decision making. Chapter 2.2 examines the emergence of the less-is-more effect in a real world prediction task. Chapter 3 is devoted to the recall principle. As a first step, in Chapter 3.1 the descriptive validity of the recall principle in a frequency judgment task is tested. Do people use it? In Chapter 3.2 I develop candidate process models based on the recall principle and test these models against an alternative inference approach (that is based on fact knowledge rather than instances). Chapter 4, finally, provides a summary of Chapters 2 and 3 and a concluding discussion, in which the findings in the reported studies are discussed in terms of their implications for the notion of ecological rationality. As the individual studies are discussed in depth in the respective chapters, this final discussion is brief and addresses the general picture only. Note that the chapters were written to be readable individually, so there will occasionally be some overlap as to definitions of terms and concepts.