Comprehensive Summaries of Uppsala Dissertations
from the Faculty of Social Sciences 131

Bounded Rationality and Exemplar Models

BY

MAGNUS PERSSON

ACTA UNIVERSITATIS UPSALIENSIS
UPPSALA 2003

Dissertation presented at Uppsala University to be publicly examined in Lärosal IV, Universitetshuset, Uppsala, Thursday, October 2, 2003 at 13:15 for the degree of Doctor of Philosophy. The examination will be conducted in English.

Abstract

Persson, M. 2003. Bounded Rationality and Exemplar Models. Acta Universitatis Upsaliensis. *Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences* 131. 48 pp. Uppsala. ISBN 91-554-5733-9

Bounded rationality is the study of how human cognition with limited capacity is adapted to handle the complex information structures in the environment. This thesis argues that in order to understand the bounded rationality of decision processes, it is necessary to develop decision theories that are computational process models based upon basic cognitive and perceptual mechanisms. The main goal of this thesis is to show that models of perceptual categorization based on the storage of exemplars and retrieval of similar exemplars whenever a new object is encountered (D. L. Medin & M. M. Schaffer, 1978), can be an important contribution to theories of decision making. Study I proposed, PROBEX (PROBabilities from Exemplars), a model for inferences from generic knowledge. It is a “lazy” algorithm that presumes no pre-computed abstractions. In a computer simulation it was found to be a powerful decision strategy, and it was possible to fit the model to human data in a psychologically plausible way. Study II was a theoretical investigation that found that PROBEX was very robust in conditions where the decision maker has very little information, and that it worked well even under the worst circumstances. Study III empirically tested if humans can learn to use exemplar based or one reason decision making strategies (G. Gigerenzer, P. Todd, & the ABC Research Group, 1999) where it is appropriate in a two-alternative choice task. Experiment 1 used cue structure and presentation format as independent variables, and participants easily used one reason strategies if the decision task presented the information as normal text. The participants were only able to use exemplars if they were presented as short strings of letters. Experiment 2 failed to accelerate learning of exemplar use during the decision phase, by prior exposure to exemplars in a similar task. In conclusion, this thesis supports that there are at least two modes of decision making, which are boundedly rational if they are used in the appropriate context. Exemplar strategies may, contrary to study II, only be used late in learning, and the conditions for learning need to be investigated further.

Keywords: PROBEX, Lazy Algorithm, Probabilistic Inference, Decision Making, Bounded Rationality, Ecological Rationality, Take The Best, exemplar models, correspondence

Magnus Persson, Department of Psychology, Box 1225, Uppsala University, SE-75142 Uppsala, Sweden

ISSN 0282-7492
ISBN 91-554-5733-9
urn:nbn:se:uu:diva-3572 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-3572)

List of Papers

This thesis is based on the following papers, which will be referred to in the text by their Roman numerals.

I. Juslin, P., Persson, M. (2002) PROBabilities from EXemplars (PROBEX): a lazy algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, 26: 563-607 (accepted)

II. Persson, M., Juslin, P. (2000) Fast and Frugal Use of Cue Direction in States of Limited Knowledge. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, : 841-846 (accepted)

III. Persson, M. Decision strategies as adaptations to cue structures. (manuscript)

CONTENTS

ABBREVIATIONS

PROBEX	Probabilities from Exemplars
TTB	Take The Best
WADD	Weighted ADDitive decision strategy

1 INTRODUCTION

Imagine yourself in a foreign pub with a lot of different brands of beer that are unknown to you. Which beer should you choose? If you have no advice to aid you in your decision you have to rely on the visual cues of the bottles and your own experience of beers. This is a hard (and important!) problem and there is no decision strategy that is certain to be a priori best.

When we compare two bottles of beer there are some different principles that could be used. Each bottle has features such as brand name, size, shape, color of bottle, kind of beer, labels, price etc. and they can all be used to make a decision. One could use a small set of important features and pick the first beer that at least fulfills one item from this list of desired features (*one reason decision making)*. A more effortful approach is to sum up all positive features for each bottle, weighted by the importance of each feature. A decision maker that uses these strategies needs to know which features are important and which are not (*weighted linear strategy*). A third strategy is to use the features of each bottle to access memories of similar past bottle of beers, which you remember as good or bad (*exemplar or case-based strategy*). Under the assumption that the features of a bottle of beer are not random but informative, it is likely that similar bottles will have similar contents. This thesis investigates under what circumstances people use decision strategies like these, for decisions that are based upon estimates, which in turn are inferred from properties (*cues*) of a decision alternative. For example, if you are asked about the population size of Heidelberg, you might know that Heidelberg has a university and some other facts that can help you guess its size.

One theme in this thesis is similarity-based reasoning that uses as many cues to make inferences. Similarity-based estimates are assumed to be inferences from similar instances in memory. If you happen to know that Regensburg has about 125 000 inhabitants and that this city is similar to Heidelberg then you can guess that Heidelberg has about the same population size. One reason decision making is a parallel theme in this thesis. It is different from similarity-based reasoning because many cues are never even considered. If you know that there is a university in Heidelberg but not in Reutlingen you can guess that Heidelberg is larger than Reutlingen, and hope that the university cue is more important than other facts you knew about these cities. One reason decision making works well when some cues are much better predictors than other cues. Under such circumstances it is possible to be both accurate and very efficient, but only if we are able to tell which cue is important and which is not. Both similarity based and one reason decision making can be seen as strategies that avoid the complexities of using a weighted linear strategy.

Cognitive psychology is the science of human information processing. Thus, our understanding of the processes of decision making cannot be complete without knowing the structure of the information (decision environments). For example, if humans adapt their strategies when they make repeated decisions and learn what works well we need to know which strategies cope well with certain kinds of information in order to develop better theories of such processes. Computational modeling of process models is an invaluable tool to explore how different decision strategies can cope with different information structures. Furthermore computer simulations provide detailed predictions from process models that would be impossible to derive analytically or by informal reflection on the models. This thesis uses computational modeling as the main tool to better understand human decision making.

The purpose of this thesis is to explore exemplar models as decision strategies in situations of uncertainty. Exemplars are stored instances of experiences that can be retrieved from memory and used as information for decisions later in life. A new theory of judgment and decision making, PROBabilities from Exemplars (PROBEX), is presented in Study I (Juslin & Persson, 2002). Furthermore, it is investigated whether PROBEX can behave as humans do in a probability judgment and decision making task. Study II (Persson & Juslin, 2000) is a theoretical investigation of the efficiency and robustness of PROBEX compared to other strategies in situations with minimal knowledge. Finally, Study III (Persson, 2003) uses PROBEX and the one reason strategy Take The Best (TTB) to predict accuracy and response times in an experiment with two very different decision environments. One where one reason decision making is enough for optimal decisions and one where exemplar based strategies is the only possible way to make optimal decisions.

2 Bounded Rationality and Decision Making

This section will present and discuss bounded rationality (Simon, 1982) as the underlying theoretical concept of this thesis. Research on bounded rationality can be seen as the study of a) cognitive processes with limited capacity and b) the information structures of the environments of the organism. The main point is that most decision making research has focused on the cognitive limitation part of bounded rationality and neglected the environment in favor of the construct of subjective utility. The implications are that the eternal question of whether human decisions are rational has been thoroughly examined on side issues as *coherence* and weak forms of *correspondence* rather than the more "down to earth" *ecological rationality*[1] (a stronger form of correspondence*)*: Does the decision process work well in the environment it was meant for?

A methodological reminder is also raised that in order to use bounded rationality as Simon intended, it is necessary to use process models of decisions processes studied in conjunction with models of the information in possible environments (*cue structures*).

2.1 Bounded Rationality

Herbert Simon developed the idea of bounded rationality in the mid-1950s, as a reaction to the use of subjective expected utility in economical theories of decision making (Augier, 2001; Simon, 1992). He thought that the mental representations and processes had to be different from the assumptions of classical rationality, where the agent knows the exact information that is needed to compute, in principle, any interesting utility. His basic idea is that the organism is adapted to an environment in the sense that, the cognitive capacity of the organism is sufficiently advanced to let the organism survive in the environment it lives in. It does not matter much if the organism violates the norms or rationality as long as it finds enough food and can

[1] Coherence and correspondence (Hammond, 2000) are defined and discussed later, but note that *weak* and *strong* correspondence are terms introduced in this thesis to make finer distinctions out of a confusing concept.

avoid danger, that is, the organism is "*satisficing"* rather than maximizing utility.

In order to develop a psychological theory of an organism that satisfices we need to know the basic perceptual and cognitive abilities the organism has and the limitations of these. In order to understand how the mechanisms can be used in a boundedly rational way, we also need to understand the nature of the environment that the organism lives in, for example how is the food distributed in the environment? Is the food easy to detect? Are there clues to where food can be found? Is it necessary to remember where the food was found? Simon likened this kind of analysis to a pair of scissors. One blade consists of the cognitive mechanisms of the organism and the other blade of how the environment is structured. Without full knowledge of both blades it is not possible to understand the behavior of the organism.

2.1.1 The neglect of bounded rationality in decision making

Simon developed his initial ideas of bounded rationality in the context of economic decision making in large organizations, and it was these ideas that gave him the Nobel prize in Economics[2], but the influence on the field of behavioral decision making has largely been indirect. Simon himself concentrated on the psychology of problem solving, and became known as one of the founding fathers of cognitive science because of his contributions.

It is interesting to read Simon's paper on bounded rationality in *Psychological Review* (1956). What he does is to mathematically model how a rat could search for food and survive in a world where food was scattered randomly. Although a thought experiment, the example given is a complete process model of the behavior of the rat, as well as a model of the environment of the rat. I believe that both the first blade of Simon's scissors, m*odeling the capacities of the organism,* and the second blade, *the analysis of the environment*, has been neglected in decision research. Rather than modelling decision processes, researchers have continued to use subjective expected utility as a norm, explaining deviations by cognitive limitations rather than providing explicit process models for the mechanisms.

Neglecting the analysis of the environment is more forgivable, because it may seem to demand divine intellectual capacity that no mortal researcher is capable of (but see, Anderson & Schooler, 1991, for an account of the relationship between memory and environment). Todd and Gigerenzer (1999) suggest that other fields such as ecology and statistics can be a source of inspiration.

[2] The Bank of Sweden Prize in Economic Sciences in Memory of Alfred Nobel is the correct name for this prize which should not be confused with the Nobel prizes in Physics, Chemistry, Medicine, Literature and Peace.

2.1.2 Two possible new directions in decision research

The most important reason why bounded rationality has not made much impact in the field (as it was originally intended) was that cognitive psychology was still in its infancy when Simon proposed bounded rationality. There was no thorough understanding of mental processes that could become the basis for a general theory of boundedly rational behavior. I will mention two areas of modern research that may contribute to such theories. The first area is theories of emotions, which cannot be neglected in decision making. The second is theories of exemplar based categorization that can be an inspiration for new theories of decision making. Emotions will be discussed briefly here and exemplar based decision making is the main theme throughout this thesis.

Emotion has often been depicted as the opposite of rational thought in philosophy (Damasio, 1998) and largely ignored during the cognitive revolution in the second half of the 20^{th} century (LeDeux, 2000). Not surprisingly emotion has been neglected in decision making as a consequence of the preoccupation with classical rationality as the measure of human reasoning ability. Within neuroscience, however, there has been a dramatic increase in our knowledge about how emotional processes is involved in memory, learning and decision making (Dolan, 2002). The *somatic marker* hypothesis of Damasio for examples states that alternative options that are reflected upon engages the emotional systems throughout the brain and body which make us feel uneasy about bad options and feel positively about good options. The slow emotional system integrates these feelings over time and helps us to choose a good alternative.

Even traditional decision research has begun to incorporate emotion in models. Decision affect theory (Mellers, 2000; Mellers, Schwartz, & Ritov, 1999) is a theory of judged pleasure in monetary gambles, where subjective expected pleasure rather than subjective expected utility determines choice.

It is thus not surprising that Hanoch (2002) has proposed that emotion is an integral part of bounded rationality. The most fascinating aspect of the latest insights of neuropsychology, is perhaps that brain damage to critical areas of emotional processing such as the orbital cortex severely handicaps the ability to handle important decisions in everyday life. Thus, it seems as if emotion is necessary for normal reasoning, contrary to what philosophers once believed.

2.1.3 Three ways of constraining bounded rationality

If we step back and examine bounded rationality in a wider perspective, a salient feature of the theory is that it is very general. If we ever meet intelligent life from other planets, it would be reasonable to develop psychological theories for such species consistent with bounded rationality. In order to formulate down to earth theories of bounded rationality, there has to be some constraints. Today there are two schools of bounded rationality that constrain their theories in an abstract fashion without touching the messy details of basic processes of perception and cognition. After a short introduction of these schools a third alternative is presented.

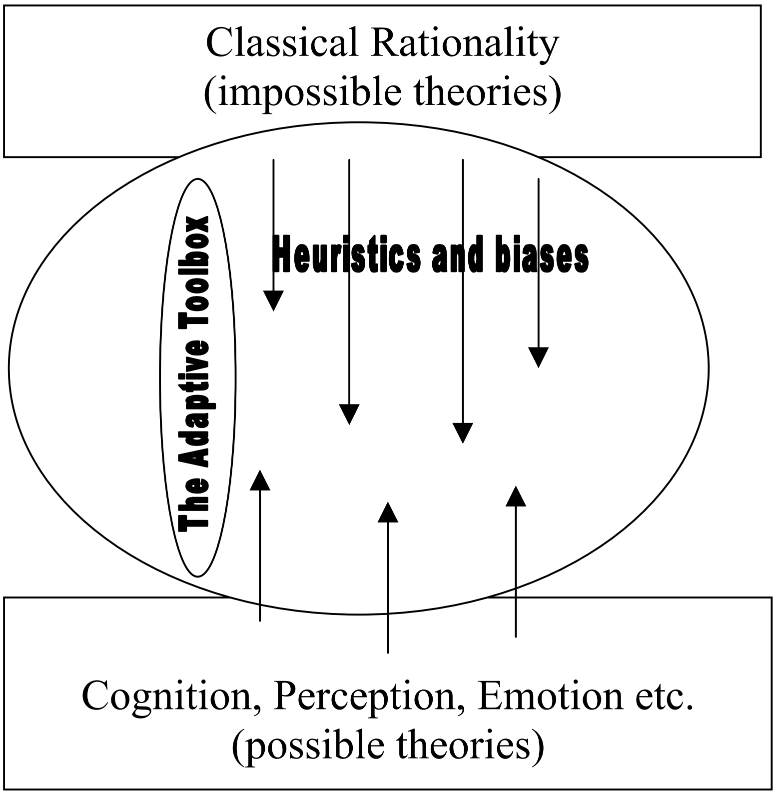

Figure 1. The big oval represents the set of all possible theories of boundedly rational decision making between empirically supported theories of cognition and the ideals of classical rationality (that cannot be true theories of high level cognition). The school of heuristics and biases has shown that people violate classical rationality in numerous empirical experiments, and has then drawn the conclusion that we use heuristics for decision making. The arrows pointing down shows that falsifying classic rationality drives this school towards cognitive theories but they have trouble going beyond vague verbal theories. The arrows pointing upwards represent the main idea behind this thesis: to simply start with process models of cognition and applying them to decision making. The thin vertical oval represents the models in the adaptive toolbox. These models have contact with cognition and are sometimes surprisingly close to optimal performance. The

problem is that the set of models that belong to the toolbox might be too narrowly defined.

The school of "*heuristics and biases*" (Tversky & Kahneman, 1974) presents itself as a theory of bounded rationality, and the constraint on theories of human decision making comes from comparing human decisions and judgments to norms of rationality. Thus, the overarching goal of this huge literature is to show that human decision making cannot be modeled descriptively with theories such as subjective expected utility and as such this research program has been very successful. This is depicted with the downward arrows in Figure 1. The connection to bounded rationality comes from the idea that we use error prone cognitive mechanisms that are simple and work well most of the time. Heuristics are mental shortcuts to easy decisions such as, for example, availability and representativeness (Tversky & Kahneman, 1974). Biases are systematic deviations in our ability to make judgments such as belief bias and overconfidence, to name a few. As a theory of bounded rationality it is very heavy on the blade of limited cognition. The blade of the environment has been replaced with rational norms applied to formal problems of logic and probability. As pointed out by Lopes (1992), proponents of the heuristics and biases approach gives only 'an honorific nod to a distinguished figure than it is an acknowledgement of significant intellectual debt' (p. 232) when Simon's views of bounded rationality is cited. The reference is made but the concept of bounded rationality in not important to the heuristics and biases school of thought.

The other school constrains its version of bounded rationality as much as possible. The "*adaptive toolbox*" (Gigerenzer & Todd, 1999) is a research program where process models are assembled from a small set of heuristic building blocks, such as heuristic principles for guiding search, stopping search and making decisions. The result is a set of "*fast and frugal heuristics"*, where each heuristic is adapted to an environment with a particular distribution of information. These fast and frugal heuristics have been compared to statistical methods such as multiple regression and Bayesian networks in computer simulations and compete surprisingly well with them (Gigerenzer & Goldstein, 1996; Gigerenzer & Todd, 1999). For some data sets they are even better than complex methods. They are fast because they require very little computation and frugal because they do not look at all the given information. In theory these new heuristics seems to be perfect candidates for a detailed theory of bounded rationality and as a side effect they revive the hope that humans are surprisingly rational after all. This school is depicted as a thin vertical oval in Figure 1, reaching from simple processes towards rationality although narrowly defined such processes.

The third alternative to constrain new theories is to look for successful theories from other fields in cognitive science such as exemplar or connectionist models (Chater, 2000; Juslin & Persson, 2002) and adapt them to decision making tasks. This approach may help us one step further towards unification of cognitive microtheories (A. Newell, 1992). This thesis concentrates on exemplar theories, but theories of emotion and any other kind of brain process could of course also become a part of this alternative vision of bounded rationality. This idea is shown with the upward arrows in Figure 1. Classical rationality as applied to higher cognitive processes has to be impossible, but this does not exclude that the brain computes optimally on the neural level. This is why the oval in Figure 1 overlaps with the box containing "impossible" theories. As long as a process can take advantage of the massively parallel computing power on the neural level, there is no upper limit to the performance of the process.

Proponents of the heuristics and biases tradition and the adaptive toolbox, would probably argue that they do hang on to the developments in cognitive theory, but there are some shortcomings that are unavoidable within each of these schools. Tversky and Kahneman (1983), for example, write of similarity based reasoning in support for the representativeness heuristic in a way which proves that they were well versed in the literature of the time, and Tversky was a leading expert in similarity judgment with the contrast model (Tversky, 1977). But there is no formal connection between cognitive models and representativeness and nor could it be, because all the heuristics of this school are verbal theories (Gigerenzer, 1996). Thus the theories of heuristics and biases deteriorate into a theory of bounded performance, rather than a theory of bounded psychological mechanisms.

The adaptive toolbox is in a much better position because the theory is based upon process models and includes references to basic processes as a prerequisite to fast and frugal heuristics. The recognition heuristic (Goldstein & Gigerenzer, 2002) utilizes the parallel processing of perceptual systems to make surprisingly accurate inferences in a variety of situations. Germans are better than Americans to judge the relative sizes of cities in USA, because they only recognize the largest cities while Americans cannot use this information since they recognize small cities as well as large cities.

But there is a fundamental problem caused by the separation of fast and frugal heuristics from the rest of cognition and perception. For example, the recognition cue is seen as a simple binary variable, which make it possible to separate the process of recognition from the simple binary decision that follows. But, there is a possibility that such decisions are not as simple. In Study I decision makers was modeled to *not know cues* as a function of how well cities were recognized. This lack of knowledge influenced the similarities between cities such that partial recognition ("I know little about this city, but a little bit more about that city") could be used to make graded similarity judgments. The decision making process then cannot be seen as independent from the recognition process. Another problem is that the decision process itself is supposed to be serial. The adaptive toolbox only includes processes that consist of a search rule, a stopping rule, and a decision rule. This is perhaps enough to study decision making in static and isolated situations. But it will be very difficult to explain the power law of automatisation (when a task is performed faster and faster if it is repeated) with fast and frugal heuristics, whereas the assumptions of parallel processing of multiple representations has been a successful way of modeling this phenomenon in categorization tasks (Logan, 1988; Nosofsky & Palmeri, 1997). In such models learning increases the amount of memory traces that can be activated which imply faster response times. Another example is Lamberts (2000) who models perceptual categorization with parallel independent feature detection processes that determine the time it takes to categorize an object.

2.2 The Role of Rationality

There are two major views of rationality that cause confusion in regard to the issue of human rationality: coherence and correspondence (Hammond, 2000).

Coherence is when the action and thoughts of an agent conforms to logic and is consistent with norms of probability and utility. Strictly speaking coherence in it self, is only a state of the organism with no contradictions. The "rationality" then has to come from some definition of optimality relating to the outside world that can tell apart good actions from bad. A paranoid person may be perfectly coherent in thoughts and actions but is nethertheless judged irrational by the society. In order to make a strictly coherent trade off between immediate rewards and learning in a changing environment one must build a complex model of the entire environment. Understanding rational behavior in terms of coherence thus rapidly becomes exceedingly complex and the decision maker has to utilize a demon of unbounded rationality (Gigerenzer & Todd, 1999) in order to compute decisions coherently.

Correspondence defines rationality as when the thoughts and actions of an agent correspond to properties of the environment or work well in a given environment. It is thus not necessary to be completely coherent as long as the results are good. I will here make a distinction between weak and strong correspondence. Weak correspondence is coherency "in disguise" where for example the decision maker is potentially rational only if a mental representation (such as a judgment of a numerical quantity) is close to the objective value in the external world. Strong correspondence on the other hand requires that the actions that follow from whatever mental representations are used work well in the given environment. For example an individual who goes to an auction with a limited amount of money, and is very uncertain about the real value of items for sale might use the simple strategy to only bid as long as the price is much lower than the uncertain estimate. With some luck there will be some item that can be bought at the low price. In most cases the "luck" was simply that the estimate was too high, but then the low price is close to the real value so it is not a loss. Occasionally this strategy would probably buy something at a really good price. This strategy is an example of strong correspondence, where the highest bids will be lower or close to the correct value despite the absence of weak correspondence.

The main point here is that it is not meaningful to ask people to give estimates of numbers they might not actually use in the task of study.

The heuristics and biases approach uses rationality defined as the coherence of beliefs and preferences in order to develop diagnostics for judgmental heuristics (Kahneman, 2000). Kahneman and Tversky (1996) also emphasize that they have studied the correspondence between for example estimates of numerical quantities with the objective numbers. But does it matter if the estimates deviate from the objective numbers, as discussed above in terms of weak correspondence? Lopes (1991) argues in other words that there is a problem of generalizability from the use of simple pencil and paper problems in the heuristics and bias tradition. It seems to be taken for granted that strong correspondence follows if and only if weak correspondence is a property of human decision making, and with that logic it is enough to demonstrate human irrationality by showing the violation of weak correspondence. The strong point of ecological rationality is that strong correspondence is possible without weak correspondence, as long as the decision process is adapted to the environment.

2.2.1 Ecological rationality

The strategies of all organisms have all evolved in highly complex environments. A psychological mechanism is *ecologically rational* if there is an environment where this mechanism is efficient and accurate. The assumption behind the adaptive toolbox is that it contains many specialized tools, and that each tool is rational in a narrow domain. If there is a tool we can use we will be rational, but if there is no tool we will not do well.

Ecological rationality is thus equivalent with strong correspondence with the addition that increasing specialization is necessary to gain high performance. A theory such as subjective expected utility is in contrast as general as a theory can be and an organism using it is inefficient (Simon, 1992). The heuristics and bias tradition clearly demonstrate that this also is the case. But since they are preoccupied by testing a theory that cannot be true, this research cannot give us any insights whether we are rational in the ecological sense. In defence of the heuristics and biases tradition it is fair to point out that the heuristics such as representativeness and availability could be examples of ecological rationality but the problem is that these heuristics are too vague to give detailed predictions and, as Gigerenzer (1996) emphatically claims, we need process models to do so.

2.3 No Process Models of Decision Making

Recently several process models of decision making have been proposed so the following critique of the field is a historical account (see Dougherty, Gettys, & Ogden, 1999, for a similar discussion).

The statement that there are no process models of decision making is provocative considering that there are plenty of simple decisions rules that each could be seen as a simple process model (Montgomery & Svenson, 1976 lists 13 different decision rules) and numerous theories of decision processes (e.g.) Beach & Mitchell, 1987; Klein, Calderwood, & Macgregor 1989; Montgomery, 1989; Svenson, 1992; Svenson, 1996). The statement is only true if "process model" is defined strictly as follows. A process model should be a) implemented as a computational model and b) is used to predict human decision behavior from c) the cue structure of the task. There are a lot of theories of decision processes, but with the given definition few of them would qualify as process models of the decision process. One exception is Huber (1994) who points out that decision making could benefit from computer simulations of process models and provides the Chunking-By-Similarity model as an example. In a review of decision making research from a process tracing perspective Svenson (1996) also mentions that cognitive modeling would be a valuable contribution to the field of decision making.

The complex developments in decision research are better explained elsewhere (Goldstein & Hogarth, 1997; Hastie, 2001; Svensson, 1996) but as a crude simplification researchers in decision making have mainly used two different methods to study decision behavior. The most common method is to collect data about the actual decisions made, and then compare them to the norms of classical rationality such as maximizing subjective utility. In this paradigm the subjective utility is either provided explicitly in the task

(Payne, Bettman, & Johnson, 1993) or are assessed from the subjects before they make the decisions (Klein, 1983).

The second and less common alternative is to directly study the decision process with process tracing methods such as verbal protocols or information boards (Svensson, 1996; Harte, Westenberg, & van Someren, 1994). This school cares less about which decisions are made, instead it asks questions about which strategies are used and how different contingencies affect the decision process. Researchers using process tracing methods are interested in decision making on a very general level with a lot of alternatives to choose from, many conflicting goals, and naturalistic decision problems. For example the Differentiation and Consolidation theory (Svenson, 1992; Svenson, 1996) has four different levels of decisions, 7 decision rules and incorporate both pre- and post decision processes. Developing a process model that can handle this complexity is not feasible. In this thesis each process model covers one decision strategy only and there are many strategies that are ignored. There are, for example, strategies that eliminate alternatives as a first step, but these strategies are left out because this thesis is limited to decisions with two alternatives, or judgments of one object.

Image theory (Beach & Mitchell, 1987) and recognition-primed decision making (Klein et al., 1989) are interesting because they seem to be consistent with exemplar models although originally intended as descriptive models of decision making. Image theory is a very general theory but if exemplar retrieval can be seen as one kind of image representation it may fit into this framework. Recognition-primed decisions theory describe situations where experts immediately retrieve a solution from memory and never consider any alternatives (Lipshitz, Klein, Orasanu, & Salas, 2001). This is clearly equivalent to what exemplar models would predict of experts.

Payne et al. (1993) are closest to the definition of process models above. They use process tracing methods and monte carlo simulations of different decision strategies, but they also stick to subjective utilities using a weighted additive decision strategy (WADD) as the optimal way of integrating cues. When they test these strategies in different task environments they manipulate the situations (for example adding more alternatives) rather than the nature of the information structure. In their experiments WADD will always give the optimal decisions as the optimal linear utility weights are given to the participants. But as is demonstrated in study III it is easy to create a task where the relationship between the given cues and what is considered the best choice is nonlinear. In principle it then follows that the cue structure of the environment is an important factor for the decision maker that cannot be ignored.

2.3.1 Two kinds of process models

Linear models have permeated judgment and decision research (see Doherty & Brehmer, 1997, for a review and discussion of the role of linear regression), because as a first approximation of any kind of data they are very general tools. As a consequence it is very difficult to show that a process model is better than a linear model (Einhorn, Kleinmuntz, & Kleinmuntz, 1979). Recently though, several new process models of decision making have been proposed: Minerva-DM (Dougherty et al., 1999) and PROBEX (Juslin & Persson, 2002) are both exemplar based models, whereas TTB (Gigerenzer & Goldstein, 1996) is based on a lexiographic rule. Minerva-DM is a memory model and as it does not use the features of objects to represent the exemplars it cannot capture the cue structures of the environments. Instead the focus is on how the assumption of multiple memory traces affects likelihood judgments. As a process model of decision making the rationale behind Minerva-DM is parallel to PROBEX but as a model of bounded rationality it fails because it neglects the impact of the environment.

Exemplar models of categorization have been very successful (see section 3). Despite this, Smith and Minda (1998) have recently advocated the old idea of prototypes as a fundamental form of representation that is more important than exemplars. New studies, however, concerning the early learning process of categorization shows that only exemplar models can explain data late in learning and that there are shifts from simpler rule based representations (rather than prototypes) towards exemplars (Johansen & Palmeri, 2002). SUSTAIN, is a neural network that builds simple representations as a first assumption and then adds more clusters of neurons as needed when errors are made. Hence, categories that can be described with rules, prototypes or exemplars, are all different aspects of the same learning and categorization process using a unitary but flexible form of representation (Love, Medin, & Gureckis, in press). It is mentioned here to show the full potential of this type of process model. Automaticity, the phenomenon that experience speeds up performance of tasks that are already learned to perfection, is also a strong point of exemplar theory. The simple explanation is that over-learning of multiple exemplars provide relevant information to automatic tasks faster, since many exemplars increases the chance that a relevant exemplar is activated early. Exemplar theory may thus potentially explain the use of simple representations in decision making, the learning process and the automatization of the decision task in a coherent theoretical framework that is implementable in a neural substrate.

Fast and frugal heuristics are examples of both bounded and ecological rationality, since they are not only fast and frugal, but also specific strategies adapted to particular kinds of ecology. A typical example of a fast and frugal decision strategy is the Take The Best (TTB) algorithm (Gigerenzer & Goldstein, 1996). TTB is an algorithm that searches for cues in a particular order and applies the first cue it can use. Figure 2 provides a simple example with four cues where it searches three cues and ignores the last cue. Take as an example the task of judging which out of two German cities has the largest population. There is a lot of cues that can differentiate these cities such as if they are in former East Germany, whether they have an international airport, a football team in Bundesliga and so on. If one of these cues is true for one city but not for the other, then the probability that the former city is the largest is greater than chance. This probability, as defined for all pairs of relevant German cities, is the cue validity for that particular cue. TTB searches the cues in the order of decreasing cue validities. That is, TTB searches the best cue first then it searches the second best and so forth until a cue is found that differentiates the cities. If for example the differentiating cue has the cue validity .7, then the decisions will be correct in 70% of those cases assuming that random pairs of German cities are chosen. Some cues, as "the city is a capital in the country", can have a perfect cue validity of 1.0, but if the cue only is true for one single city, it cannot be used very often, because the cue is false for both cities in almost every possible pair.

TTB is fast because it only spends time on retrieving and comparing a few cues for each decision. It is very rare that it has to search all cues. TTB is also frugal in the sense that it does not assume any complex computations or complex representation in the moment of decision making. A problem though, as will be discussed later, is where the cue validities come from, because even if the definition of cue validity is computationally simple is it not as fast and frugal as the decision part of the algorithm. Despite the hinted drawbacks though, TTB does have a solid merit. It can compete evenly with complex algorithms such as multiple linear regression and even Bayesian networks in terms of accuracy (Gigerenzer & Goldstein, 1996; Martignon & Laskey, 1999).

In summary both kinds of process models are good candidates of bounded rationality. The main difference may be that exemplar strategies are more general strategies while TTB is a more specialized strategy. Control of the cue structures then may be crucial, to see which of these decision strategies that best explain human decision making. It also important to differentiate these modes of decision making from linear weighting strategies that may confound the results.

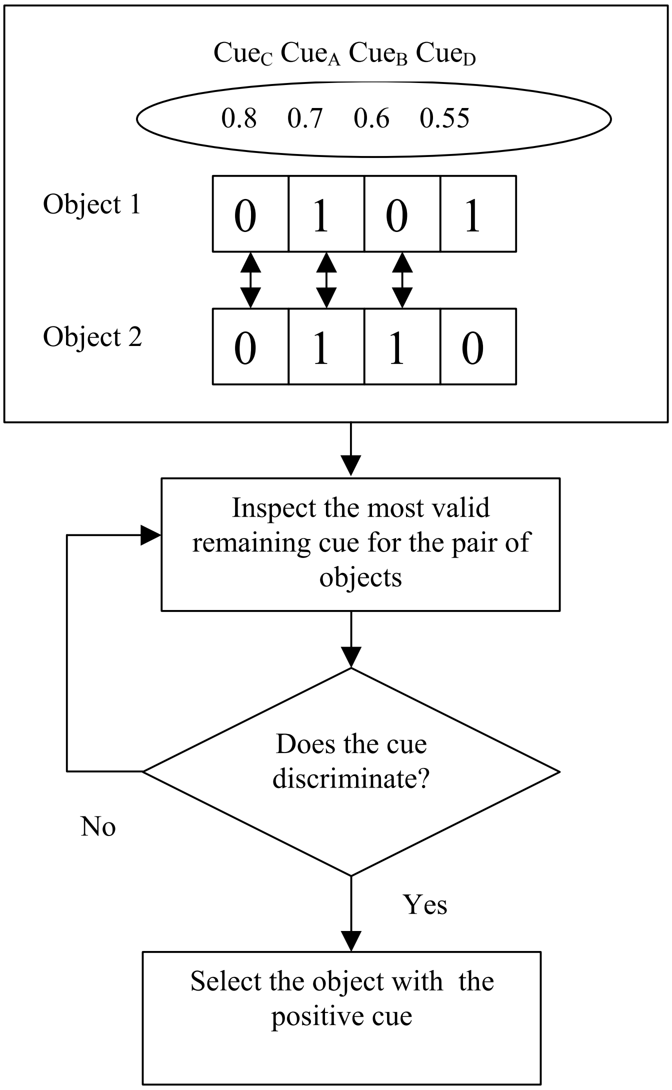

Figure 2. Take The Best compares two objects with binary cues and make a judgment of which of these object has the highest value on an unknown feature, the criterion. In this example there are four cues A, B, C and D. TTB searches these in the "best" order. It first compares cue C, which has the highest cue validity. The cue validity is the probability that a correct choice can be based upon the cue on its own. Since both cues C and A are similar for the object 1 and 2, a decision is not made until cue B is compared. Note that cue D is ignored. In most cases TTB will make the decision searching only the first or second cue and ignore all other cues.

2.4 General Aim

The aim of this thesis is to develop and empirically test a new theory of decision making. It is based upon exemplar models of categorization, which at the computational level is different from what has been common in decision making. A secondary aim is to compare this new theory with the ideas of “The Adaptive Toolbox” (Gigerenzer & Todd, 1999). From the point of bounded rationality it is expected that human decision making is ecologically rational, and uses several decision strategies that suits different situations. The new theory is presented next and then the studies follow.

3 PROBEX

PROBEX is based on the generalized context model[3] (Nosofsky, 1986), and also incorporates some stochastic noise components from the combined error model (Juslin, Olsson, & Björkman, 1997; Juslin, Wennerholm, & Winman, 1999). It is both a model of inference making and the subjective probability that each elicited inference is correct. It has no explicit bias parameters. Any kind of bias has to come from the structure of the model itself, in interaction with the information that is given.

The first exemplar model was the context model proposed by Medin and Schaffer (1978), later extended into the Generalized Context Model by Nosofsky (1986). Exemplar models have been applied successfully to a range of different tasks such as, for example, memory (Hintzman, 1988), attention (Logan, 2002), categorization (Lamberts, 2000; Medin & Schaffer, 1978; Nosofsky, 1986; Nosofsky and Palmeri, 1997; Kruschke, 1992; Storms, De Boeck, & Ruts, 2000), automatisation (Logan, 1988), social cognition (Smith & Zarate, 1992) and language (Daelemans, 1995).

The principles behind PROBEX are not new. In artificial intelligence and similar disciplines it would belong to a class of algorithms using *memory based reasoning.* These algorithms are based on stored experiences or any kind of raw data. Examples of the tasks involved range from classification of news articles to robotics (Kasif et al., 1998). *Lazy algorithms* as defined by Aha (1997), is a term that captures memory-based reasoning that fulfils the following criteria for lazy information processing: First, stored data is not processed until a request for information is received. Next, the result of the request is based on a combination of stored data. Finally, the request and the result are discarded. PROBEX is a true Lazy algorithm since it only uses stored training exemplars for inferences. It uses locally weighted regression as the method of inference, but since it is a cognitive model it also has stochastic elements that capture the random aspects of any psychologically plausible decision process.

The purpose of PROBEX is to show that lazy algorithms in general and exemplar-based models in particular can be good examples of bounded

[3] The most important difference to GCM is that PROBEX does not use attention weights on each separate dimension. Instead it only has one similarity parameter *s*.

rationality as discussed above. Furthermore PROBEX is supposed to show ecological rationality in competition with other algorithms as in Study I and II, and match human performance in their ability to make decisions, estimates, and subjective probability judgements, as presented in Study I. An empirical test of PROBEX in a decision task with learning is presented in Study III.

3.1 The Model

A formal presentation of the model is given in Study I and II, and the following description should be enough to understand the principles behind the model. Figure 3 provides a schematic overview of the model without going into details.

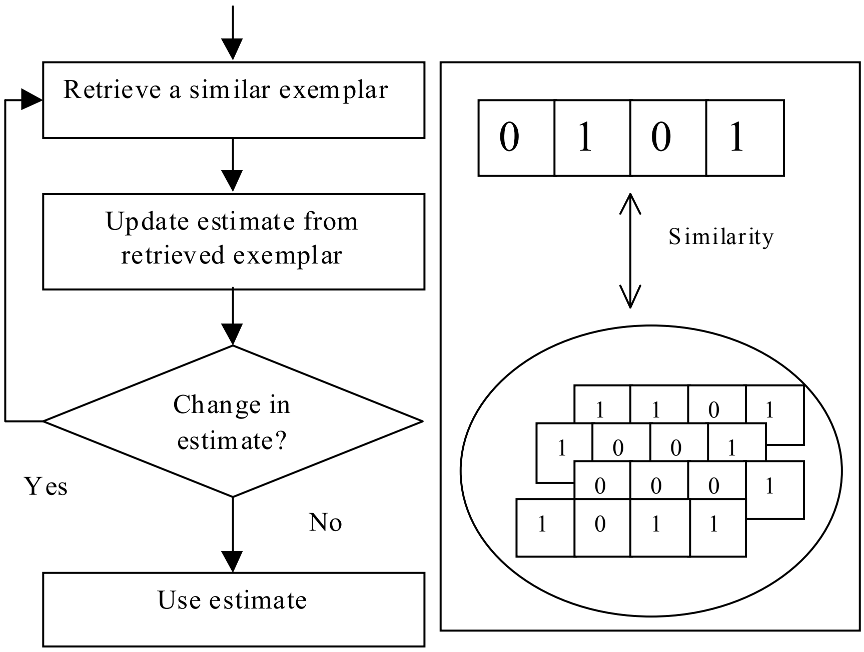

Figure 3. A simplified overview of PROBEX. To the right an object with a number of features is presented and this object is compared to all stored exemplars in memory. One of the most similar exemplars is activated and retrieved from memory. Any useful information retrieved is used to make an estimate for the object. This procedure is repeated until the estimate do not change, and then the process terminates.

The exemplars are represented as vectors of numbers, or as points in a multidimensional psychological feature space. One dimension or feature of these vectors, the *criterion*, is always special in all the examples given in this thesis. The criterion is the feature the algorithms are supposed to infer. The inference task can be to estimate the criterion for a new object, or to estimate this value when it is unknown for a stored exemplar. The model could in principle make inferences for any feature, but for simplicity and clarity the distinction between criterion and features is made. Given a task concerning an object with an unknown or hidden criterion, the model makes an inference of the most likely value, based on the exemplars in long term memory.

The estimation process is based on serial sampling of exemplars where the order is probabilistically decided. The sampling probability for any exemplar is proportional to the similarity between that exemplar and the object in question. The similarity is a non-linear function of the city block distance metric between the compared exemplars in multi-dimensional space, and for binary features it can be easily computed as the similarity constant s raised to the power of the number of differing cues.

Every time an exemplar is sampled a tentative prediction is made, calculated as the average of the criterion of all sampled exemplars so far, weighted with the similarities of those exemplars. Without this weighting, the estimate would simply become the mean of those exemplars that are sampled, no matter which features the object has.

Any serial process needs a stopping rule and in this model the rule is rather simple. If the relative change in two successive tentative estimates do not change more than a certain proportion k of the current estimate the process stops and the last estimate is elicited as the overt response. The principle behind the rule is to avoid retrieving more exemplars than necessary, without using any complex statistical analysis of the sampling process.

If the task is to assess the probability that a statement is true, such as estimating if the value of a painting is higher than the given price, then the process is the same, in the sense that the sampling procedure is similar. The difference is that the similarities are summed for each exemplar sampled that satisfies (value>price). When the process terminates as detailed above, the overt subjective probability response is calculated as the final sum of such similarities divided by the total sum of all sampled exemplar similarities.

The response time is modeled as proportional to the number of samples.

3.2 The Parameters

As with most process models some quantitative aspects are unknown and four parameters are necessary in PROBEX. There are two for the main estimation process and two more for the subjective estimate of probability.

The *similarity parameter s* decides the degree of nonlinearity of the similarity measure. It is in the range from 0 to 1. Small values of s means that only the most similar exemplars have a high probability of being sampled and since the weights in the calculations are the similarities, very dissimilar exemplars will not contribute much to the estimates even if they are sampled. A high value of s makes the algorithm sloppy in the sense that all exemplars are perceived as similar to the object in question.

The stopping rule is governed by the threshold of the minimum allowed change in the estimate k. High k means speedy inferences with few sampled exemplars and k close to 0 means that PROBEX carefully samples most or even all exemplars.

The subjective probability estimate needs to incorporate the fact that nothing can be judged with certainty if none or few exemplars are sampled. Thus a dampening mechanism controlled by the parameter ϕ is added to the ratio between summed positive similarities and the sum of all sampled similarities (Andersson, 1990; Nosofsky, Kruschke, & McKinley, 1992). In the case where no exemplars are sampled this addition to the equation makes the subjective probability 0.5 default for pure guessing. Without this modification the model would predict severe overconfidence with increasing ignorance. Nobody is perfect and a normally distributed response error (Juslin et al., 1997; Juslin et al., 1999) is added to the subjective probability as controlled by a response error variance parameter. The resulting probability is truncated to the nearest valid response if it is below 0 or above 1.

3.3 What is new? Assumptions and Additions

PROBEX is not a model of perceptual categorization and classification (Lamberts, 2000; Nosofsky & Palmeri, 1997; Palmeri & Flanery, 2002), and differs from such models in three respects. First, it does not assume learned categories, thus it can model inferences from the learned exemplars without previous training.

Secondly, it deals with general knowledge tasks, where semantic knowledge is accessed in memory, rather than learned representations of perceptual stimuli. The most important implication of this is that the timescale of the response time predictions is likely to be different compared to models for perceptual categorization, since perceptual activation of exemplars is much faster than semantic activation. The tip of the tongue phenomenon is an extreme example of this, where recognition is virtually instant, but recall may take an eternity. The sampling process is assumed serial and not parallel which is quite unorthodox for an exemplar model.but the Exemplar Based Random Walk (EBRW) model (Nosofsky & Palmeri, 1997) also uses serial sampling.

The third distinguishing property of PROBEX is that it is not assumed that multiple exemplars of the same external object are stored. Thus the criticism against the idea of exemplars as a completely unrealistic waste of memory does not hold for PROBEX. The representations needed for PROBEX are simply those memory fragments that happen to be stored in long term memory; a very attractive property of a model of bounded rationality. In sum: Neither complete encoding and storage, nor exhaustive retrieval is presumed by PROBEX. Most, if not all, of the assumptions of PROBEX are the same as those made by Smith and Zarate (1992) who extended exemplar theory to social judgments.

3.4 Psychological Plausibility

Gigerenzer and Todd (1999) have argued that their algorithms are frugal in the sense that they do not need any complex computations at the moment of decisions. Compared to PROBEX this is a good point, but there is another side of the coin: the need for precomputed abstractions.

First, algorithms such as Take The Best must know the best order that the cues should be searched, which means that some sort of calculations or inference has to be made in advance. PROBEX, on the other hand, does not need any representations computed in advance, it just stores exemplars. Second, those algorithms that need some special representation to work properly has to know the task in advance, or will be forced to compute those representations from memory in novel situations. As emphasized in Study II, the third advantage of PROBEX is that the learning and use of cue directions is an automatic benefit of the similarity comparison process. A cue sorting algorithm like Take The Best, does not integrate scarce information well and need more exemplars than PROBEX in order to work well.

The arguments above derive from the fact that PROBEX is a Lazy algorithm. As such, PROBEX is rather unsophisticated compared to state of the art Lazy algorithms in the field of artificial intelligence. The main strength of PROBEX is that it captures the essentials of what might be psychologically plausible, within the memory learning paradigm and bounded rationality.

3.5 Serial vs. Parallel

Implemented on a serial computer PROBEX is very slow, because it has to serially simulate the properties of a parallel associative long term memory. The sampling process though is assumed to be serial and in terms of speed PROBEX operates with the basic units of time it takes to retrieve each exemplar. Take The Best scans cues of one or two exemplars in a particular order that has to be fetched from memory. It is hard to compare the timing of these different kinds of mental operations, as it is like comparing apples with oranges.

3.6 Nonlinearity

An area where PROBEX clearly shines is when it is applied to non-linear data, because exemplars models can in principle model any kind of function (Rachlin, Kasif, Salzberg, & Aha, 1994) given an infinite source of exemplars and a suitable similarity parameter. This theoretical state of perfect knowledge is of course not within the reach of PROBEX which aims at psychological plausibility. But it can learn exceptions to otherwise linear data, and it can handle the classical "exclusive or"-problem (Klemm, Bornholdt, & Schuster, 2000). Linear models and cue sorting algorithms cannot manage these problems without some transformation of the cue patterns, which may in the end amount to using an exemplar based mapping of patterns!

4 STUDIES

4.1 Study I

The purpose of this study was to investigate the ecological rationality of the new model PROBEX with computer simulations, and then to assess its psychological plausibility by comparison to empirical data.

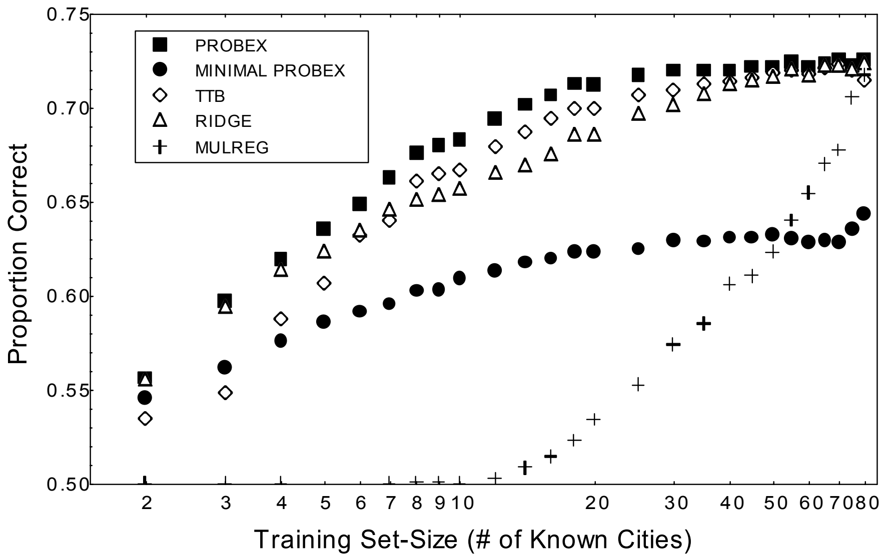

Figure 4. Simulation results from the German city population task. Minimalist PROBEX is a version of PROBEX that samples only one exemplar. RIDGE and MULREG are two types of multiple linear regression.

PROBEX is compared to other decision algorithms in the German city population task, or simply put: "Which out of city A and B is the largest?" The problem is how to make the best guess given a random amount of information about the cities in the country. The data set used is the 83 largest cities in Germany, which was the original data set first studied by Gigerenzer and Goldstein (1996). The given cues (9 in all) are only loosely connected to the population of the cities such as “Does the city have a university?”, or “Does it have a football team in Bundesliga?” and so on. The chosen cues should capture some everyday realistic facts that can safely be assumed to be available for arbitrary general knowledge estimation tasks.

4.1.1 Simulation results: the ecological rationality of PROBEX

In order to test the algorithms as severely as possible, the data set with German cities was divided randomly into a training set with known city sizes and a test set with no known sizes. The algorithms were allowed to use the training set as exemplars, or as data to calculate optimal weights, cue validities, correlations or whatever was required. Every possible pair of cities from the remaining test set was used to test the algorithms on the pair-comparison task. The measure of accuracy was the proportion of correct answers as a function of the number of cities in the training set.

The results are shown in Figure 4 and it is easy to see that for each training set size PROBEX achieves better than or equally to all other algorithms. There is no need for significance testing because each data point in the graph is based on 1000 randomly drawn training sets, which is enough to eliminate the variance of the estimated means to what is barely visible.

The pair-comparison task was used to compare PROBEX with Take The Best which cannot make predictions directly about the criterion (population). Further simulations examined algorithms that make quantitative predictions and the results are shown in Figure 5 where the mean absolute deviation of the predicted city sizes is plotted for each training set size and algorithm. These results are not as clear cut as in Figure 4 but overall it is clear that PROBEX is the “winner” again.

In the competitive simulations described above all algorithms were allowed to use all information in the training set which is not plausible. In the next simulation the accuracy of PROBEX was explored with different parameters (see study one for details). The main results is that if the stop rule is disabled with $k=0$ then the similarity parameter s is not important. But if only a few exemplars are sampled, a low value for s is important in order to maximize the chance that relevant exemplars are indeed sampled. Is it then possible to always use a very small value of the s parameter? Perhaps, but there is a limit to how small s can become because numerically computed similarities will become almost infinitely close to zero for decreasing values of s, and if some neural noise is added then the ability to discriminate between exemplars is lost. Thus, s has to be chosen such that the magnitude of the similarities is larger than the noise in the system.

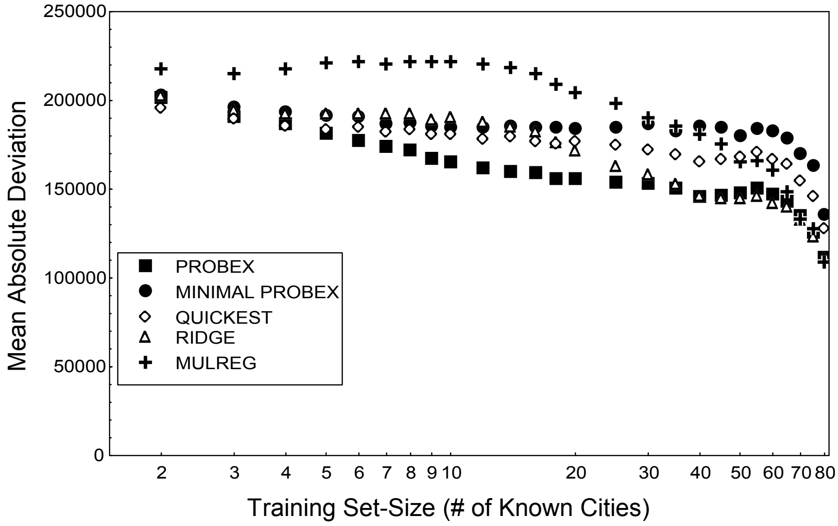

Figure 5. Mean absolute deviation of predicted city sizes as a function training set size. Low numbers are better. QUICKEST is an estimation algorithm similar to Take The Best. It assumes that the city to be judged is small, but increases this estimate for every positive cue it finds and stops as soon as a negative cue is found.

In conclusion it has been shown that the model is robust and can compete in accuracy with a wide range of algorithms and the similarity driven inference does not need a very specific or clever selection of parameters in order to work. A parameter combination of a high value of *k* and a robust low *s* should be fast and frugal and still fairly accurate.

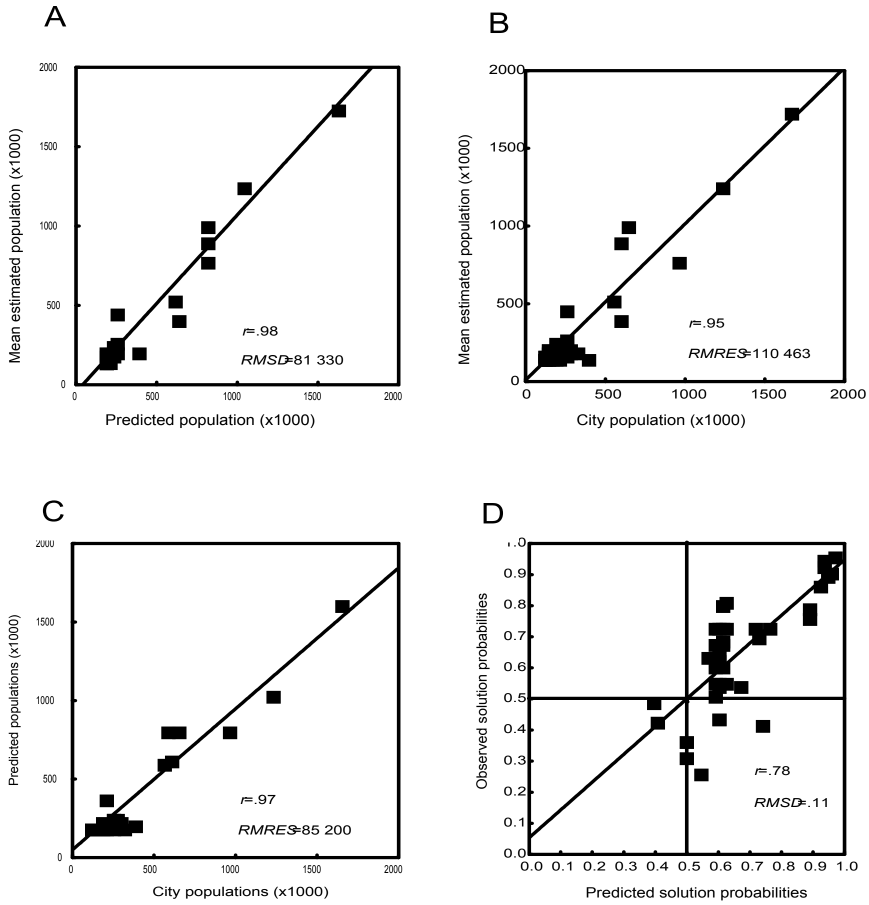

Figure 6. (A) Mean point estimates of German city-populations by 40 particpants plotted against the predictions by PROBEX. (B) Participant's mean estimates plotted against the true values (human achievement). (C) PROBEX' estimates plotted against the true values (PROBEX achievement). (D) The observed solution probabilities plotted against the predictions by PROBEX.

4.1.2 Empirical results: the psychological validity of PROBEX

The simulations showed that the algorithm is indeed very good in principle. The remaining question is if the model has the same properties as human performance. As a test of human performance, 40 participants were given the hard task of guessing the size of 40 German cities and provide confidence judgments of these estimates in three different response formats: half-range, full-range and interval format (Juslin et al., 1999). Confidence judgments are only mentioned in this summary when necessary, but see study I for definitions and results.

Then the parameters of PROBEX were fitted to the human data. The procedure was to minimize the sum of the normalized root mean square deviations (RMSD) of all dependent variables. There were 8 dependent variables, 4 for the full-range and the same 4 for the half-range condition. Further details are given in the Appendix of Study I. The dependent variables were the point estimates of each city, the proportion correct for each city, the calibration curve and the confidence judgment distribution.

In terms of what humans seems to do, assuming PROBEX is a proper model thereof, the fitted parameters can be described verbally as follows: We use as few exemplars as possible and we are picky in the sense that only the most similar exemplars are retrieved. These are the parameters that define "fast and frugal" use. Our ability to make subjective probability judgments are influenced a lot by the response error and the dampening factor plays a large role because little information is provided about the criterion when few exemplars are retrieved.

How well does the model fit data? Figure 6 give an overview for the population estimates. There are 3 possible ways to plot the population estimates of the participants, predicted population estimates of the model and the real populations of the German cities. These plots are shown in Figure 6A, 6B and 6C. The data for the half-range and full-range format conditions has been merged into one data set in this figure, since the estimates do not differ much depending on the response format of the subjective probability responses. All the correlations are very high and there is nothing peculiar to report. The model work sufficiently.

The plot between the predicted and observed solution probabilities (proportion correct) in Figure 6D is more interesting. The question is whether the model makes the same mistakes as humans. All the points in the lower half of the plot are misleading items to the participants since they on average guesses the size relative the criterion worse than chance. The model is fooled by half of these items as a result of the cue structure of the data set. The surprising thing is that there is no city that PROBEX predicts as misleading which people do not also predict as misleading (the upper left quadrant is empty). The model seems to capture some misleading aspects of the cities but not everything.

The empirical results are promising and it is safe to conclude that the model can predict behavior that capture the main features of human behavior even without a detailed modeling of the specific cues that people use (i.e. there is no guarantee that the participants use the exact 9 cues used in the simulation).

4.1.3 Conclusions

PROBEX is accurate and robust, and as discussed in the presentation of the model it is efficient with the assumptions of similarity based parallel activation of exemplars. The parameter fit of the model to data shows that two sampled exemplars are enough to reach the accuracy of the participants and to capture the main structure of all dependent variables. The parameters are psychologically plausible as well as ecologically rational. As a model of bounded rationality it has passed a first preliminary test.

4.2 Study II

PROBEX outperformed the other algorithms in the traditional German city task in Study I. The enquiry here is expanded beyond robustness, efficiency and accuracy, in order to assess what happens when there is *very little known about the environment*? First it is investigated what an algorithm has to do to perform well under such circumstances and which algorithms do that. Second, the algorithms are tested with a cue structure where algorithms that perform well with little information often fail when much is known. Finally, a simple demonstration with an evolutionary perspective shows why it is important to perform well with little information.

A good decision strategy should be useful in states of limited knowledge, that is, when there is perhaps as few as only two exemplars available to generalize from. One major problem in prediction from very few exemplars is to assess the direction of each cue. Are high values of the cues correlated with high or low values of the target variable? The ability to detect the directions of cues was tested more thoroughly in this study, by adding two algorithms based on Dawes Rule that illustrate the importance of knowing the directions of cues. Dawes rule is very simple: count the number of cues in each exemplar that has a positive correlation with the criterion and pick the exemplar with the highest number as having the largest value on the criterion dimension. The first version of the rule, named A Priori Dawes rule, does not estimate the directions of the cues but is given the true directions a priori. This version is given as a reference to what can be achieved in theory by an algorithm that only uses the directions of the cues. The second version has to calculate the directions or guess them from ordinary correlations between each cue and the target variable.

When there are only two exemplars Dawes Rule is optimal, since the cue direction is the only information given by two exemplars. As an example, if X_{big}=[1010] and Y_{small}=[0110], then we have a best guess of the cue direction for cue number one and two, but for the third and fourth the cues are the same and contain no information. It is not possible to tell which cue directions are stronger with only two exemplars, but the cue directions can be expected to be correct more often than chance predicts if the cues and the criterion are correlated.

4.2.1 Simulation 1: accuracy and cue direction

This was a replication of the simulation in Study I with some additional algorithms, and the results did not change the conclusions drawn in Study I. The interesting thing to note is that the Godlike perspective on cue directions implemented by A Priori Dawes rule is very effective in this task. It even performs better than the best asymptotic levels of the other algorithms. This is probably because it never overfits parameters from the training set that later do not generalize to the test set. PROBEX, ridge regression and Dawes rule perform optimally with two exemplars in the training set, simply because they do integrate all the scarce information about cue directions without losses. Ordering the cues, as Take The Best and QUICKEST do, and then not using all of them does not work well.

The success of PROBEX is a consequence of the fact that PROBEX and Dawes rule is equivalent in the differentiation task. When the estimates are computed with two exemplars in the training set, it is easy to show that the rank order of the predicted populations of the test set exemplars necessarily is the same for PROBEX as those given by Dawes rule.

The results confirm that knowing the directions of cues in this task is very important and that it is only PROBEX, Dawes rule and ridge regression that do this for two known exemplars.

4.2.2 Simulation 2: cue direction in a noncompensatory cue structure

Is the cue direction always as important as it seems to be in the German city population task? In order to test this an artificial data set was created. A linear structure was chosen where the optimal weights have the same structure as that of binary numbers. The advantage of this noncompensatory (each cue is always worth more than all cues with lesser weights combined) cue structure is that it allows Take The Best to asymptotically reach 100% performance. This data set can test the limitations of any algorithm that relies on detecting cue directions only, and also put PROBEX to the hardest test imaginable because it should not be able to handle linear data well with few training instances since it has no explicit mechanism for extrapolation.

Linear models fit linear data perfectly if they are given enough linearly independent data points.

The results shown in Figure 2 in Study II, shows that the regression models easily reach 100% performance, but with only two exemplars PROBEX is of course optimal. Take The Best does reach asymptotic performance, but not very fast. Further it needs 7 exemplars to beat PROBEX and is never as superior as PROBEX is for few training exemplars. A Priori Dawes rule performs a little better in absolute numbers than in the German City task, but here it has reached some kind of theoretical ceiling of about 77.5% because it cannot differentiate between the importance of cues.

The implications of this simulation is that knowing the direction of cues is the single most important factor up to 6 exemplars even in this test. Further PROBEX does manage to compete with the traditional competitors from the adaptive toolbox even on their own playing field, because it is not limited to detecting the directions of cues as Dawes rule is.

4.2.3 Simulation 3: evolution and robustness

This was a rather crude but still illustrative computer simulation of why good performance with little experience might be crucially important from an evolutionary point of view. The idea is that if you die as a beginner you will never become an expert. In Figure 3 of Study II the population development is shown for two environments with two competing species in each. Early Learner is the winning species in both cases. It has an accuracy development as a function of increasing experience similar to that of PROBEX. It wins even in the case where the competing species Late Fast Learner on average is as good as Early Learner, because the population loss in the first half of every generation is too large to be compensated for later.

The conclusion is that as a genetic adaptation over long term, PROBEX is a more likely candidate to survive. Still the term adaptive may imply that a decision strategy is an adaptation acquired during a lifetime rather than an inborn ability. The question that follows is then which strategies can be learned easily by the generic learning mechanisms given by our genes. Thus, PROBEX does not need to have a genetic origin. If humans are born with the ability to choose between similarity based reasoning and cue based search as TTB then it is likely that similarity based reasoning is chosen because it gives fairly good results early in most of the new tasks a human faces during development. Thus, the evolution of these strategies does not have to take place over hundreds of generations, but over hundreds of new tasks in the environment.

4.2.4 Conclusions

The findings here support the idea that PROBEX is a very flexible all purpose tool that performs well even in unfavorable environments, and with very little knowledge. It was also shown that if a selection of strategies is based on beginner performance then PROBEX could be favored even in circumstances where it is not the best choice for expert performance.

4.3 Study III

Distinguishing similarity-based processes from rule-based processes is difficult. Furthermore there are good reasons to believe that humans are capable of both kinds of reasoning (Hahn & Chater, 1998; Sloman & Rips, 1998). The experiments in Study III aimed to separate these processes by manipulating the *cue structure* of environments provided in a laboratory experiment with artificial stimuli. The cue structures were designed to favor two process models of decision making: Take The Best and SIMPLEX[4], where SIMPLEX is a simple nearest neighbor version of PROBEX. These processes in conjunction with the cue structures are sufficient to predict response times and thus corroborate the use of one reason decision making (TTB) and exemplar-based decision making (SIMPLEX).

There has been little evidence for Take The Best so far. Bröder (2000) found that TTB was used more when the cost of examining cues was increased, but the rate of use was nonetheless limited. These results have been replicated with a process monitoring approach (B. Newell & Shanks, 2003) using six or two cues (B. Newell, Weston, & Shanks, 2003) and it was again found that only a minority of the participants was consistent with the predictions of TTB. Juslin, Jones, Olsson and Winman (in press) investigated a multiple cue categorization task (could also be seen as a decision making task) and found little support for a strategy similar to TTB. One aim of this study was to test if TTB could be used at all.

In Study I and II the given stimuli were names of cities and the presentation of a stimuli was assumed to initiate covert processes using representations acquired from real life experience. It is close to impossible to test specific predictions from process models under such circumstances. Study III was thus designed as a repeated deterministic decision task with feedback on the criterion, where participants started with no knowledge.

The computerized task in Study III was to sell vacuum cleaners as a traveling salesman. Participants had to select one out of two alternative cities to travel to every day. Information for the decision was four binary cues

[4] Note that there is an unrelated linear programming algorithm also called SIMPLEX.

given for both cities. Feedback on the criterion was the number of vacuum cleaners sold in each city. The cue labels was randomized for each participant to eliminate any effect of prior real world experience. The participants were presented with 240 trials (or 224 trials, see the Procedure section of Experiment 1 in Study III for an explanation) of this pair-comparison task. There were 16 different cue patterns (cities) and the criterion associated with each cue pattern defined the cue structure. The dependent variables were the decisions made on each trial as well as the response time. All pairs of cue patterns were presented twice, once in the first half (Block 1) and once in the second half (Block 2).

4.3.1 Experiment 1: Cue structure and cue presentation format

If decision makers are truly adaptive they should quickly learn any cue structure and show signs typical for the appropriate process. Two cue structures were designed to favor TTB and SIMPLEX respectively. TTB belongs to the class of lexicographical strategies. Such strategies are used when a list of words is sorted in alphabetical order and to quickly pick out the largest of several numbers. With binary cues it is natural to use binary numbers such that for example “1111” is worth 8+4+2+1= 15, where 15 is the criterion. This cue structure (TTB-Friendly) is noncompensatory, that is, if a cue is true then the cues less important than this cue cannot be worth more. 1 is less than 2, 2+1 is less than 4, and 4+2+1 is less than 8. Participants need only discover the correct order to compare the cues for the alternatives in a pair-comparison task in order to be 100% accurate with this cue structure. TTB does not make estimates of the criteria of the alternatives, it simply makes an inference about which alternative is most likely to have the largest criterion. SIMPLEX on the other hand estimates each alternative independently by retrieving the criterion from a single exemplar stored in memory that is most similar to the alternative, and decides on the alternative with the highest estimate. Using memory like this allows SIMPLEX to handle difficult cue structures, such as those with a nonlinear relationship between cues and criterion. The EX-Friendly cue structure was designed to be “unfriendly” against linear strategies such as TTB and any other linear strategy but friendly to SIMPLEX. The 16 cue patterns were assigned to 8 arbitrary criterion values. The important point to note was that the pair of cue patterns attached to each criterion was the inverse of each other. Such pairs could be “1100, 0011”, “0100, 1011”, “0000, 1111” etc. TTB fails miserably with this cue structure since it can only make correct decisions for one pattern in each pair. Consistent use of TTB with this cue structure predicts a proportion correct of .5 over all pairs of cities in the experiment.

Which sequence of binary cues is easiest to remember, “GHHG” or “Cheap beer/No Airport/Poor cultural life/Bicycle-friendly”? These cue presentation formats were used in Experiment 1 and were named Letter and Text respectively. The former is abstract and hard to assign a meaning, yet it is very compact and easily encoded. The latter sequence may help your imagination to visualize a city, but processing the information is more difficult. It is not easy to see that the first and last cues are both positive while the second and third cues are negative. An exemplar strategy that relies on holistic processing of all cues may depend upon the format of the cues. Most support for exemplar based models comes from perceptual categorization (Lamberts, 2000; Nosofsky, 1987; Palmeri & Flanery, 2002), and it is possible that there will be no exemplar effects if cues cannot easily be processed holistically with the perceptual system.

A strong interpretation of bounded rationality is that humans have the resources to handle any cue structure given a reasonable amount of experience with it. Will participants adapt to the two cue structures: the noncompensatory TTB-Friendly and nonlinear EX-Friendly? But if they do, will they then use the appropriate strategies TTB and SIMPLEX? In the case of the EX-Friendly condition good performance is enough to assume an exemplar based strategy since no other strategy will do well. It is difficult to draw any conclusions from accuracy in a task using The TTB-Friendly cue structure because any linear weighted and exemplar strategy can be accurate with this cue structure. The solution to this dilemma is that TTB predicts large differences in response times for different trials. A linear weighted cue strategy or exemplar strategy will integrate all cues with few systematic differences in response times. TTB searches cues one at a time and stops at the first cue that allows a decision, which predicts a linear relationship between the number of cues searched and response time. SIMPLEX allows for the weaker prediction that decisions where one or both alternatives is recognized should be slightly faster as the criterion can be retrieved directly from memory.

A two by two factorial design was used with 10 participants in four groups: TTB-Friendly/Text, TTB-Friendly/Letter, EX-Friendly/Text, and EX-Friendly/Letter.

Accuracy

Performance in Block 2 (the last half of all trials) ranged from excellent to almost random between the groups. The accuracy in the TTB-Friendly/Text group was close to perfect, while the EX-Friendly/Text group was just slightly better than random. The logic of the experiment seems to rule out exemplar based decision making when the cues of the alternatives are presented as text. In this light the results from the groups where the Letter cue format was used almost turned the conclusions upside down. Participants in the TTB-Friendly/Letter and the EX-Friendly/Letter groups spread out through the whole spectrum of random to perfect accuracy. On average both groups ended up halfway between random and perfect accuracy. Assuming that the participants used TTB and exemplars as expected the conclusions is that TTB is easy to use with the Text format, but is harder to use with Letter format. Exemplars could be used with the Letter format but otherwise not.

Model fits of TTB and SIMPLEX

TTB and SIMPLEX have no free parameters, but they do have a *parameter space*. The strategies use a representation of the task: TTB needs a cue order to search the cues and SIMPLEX needs a set of learned exemplars and their criteria. The models can be fitted to data by selecting the representation that best explains data. If the participants reach optimal performance it could be assumed that they used either the optimal cue order, or learned all 16 exemplars. The problem is that both models would predict optimal performance and hence accuracy alone would not be able to tell the models apart. Another complication is that many participants did not do very well and we want to know if this is because they used the wrong strategy or simply were confused.

The models were fitted to individual participants using decision data from Block 2. That is, the models had to make decisions for every trial as the participants did. Every possible representation for the respective model was tried and the representation that best duplicated the decisions of the participant was selected. The fitting procedure of SIMPLEX was complicated because it was found that it overfitted the data such that it included many more exemplars than necessary. This was corrected by also fitting the proportion correct answers. One might think that fitting decisions and proportion correct should be equivalent, but this is not true, since the fitting procedure can overfit to those responses that were more or less random. In the case of SIMPLEX this systematically seems to add many wrong exemplars to the final fitted representation.

The individually fitted models were then used to predict systematic patterns in the response time data. If it is possible to explain some of the variation in the response times using the representations deduced from the decisions then the models have at least explanatory validity to some degree. Fitting the models to response time data directly would be harder to interpret since these models would probably fit *any* data to some extent. Model fitting has recently been criticized (Roberts & Pashler, 2000) and the solution proposed by Pitt, Myung and Zhang (2002) is very complicated.

The strong point of this procedure was that the response times were not included in the fitting procedure, and the conclusions were made from how well the fitted models could predict the response times.

TTB predicts that a decision is fast when it is made from the best cue (the first cue in the cue order) and that decisions spend increasingly more time when more cues are searched. This linear relationship was found with the TTB-Friendly/Text condition (see Figure 7) and for the TTB-Friendly/Letter as well. The only deviation is that on those rare occasions where a decision has to be made from the fourth cue decisions seem to be faster than expected.

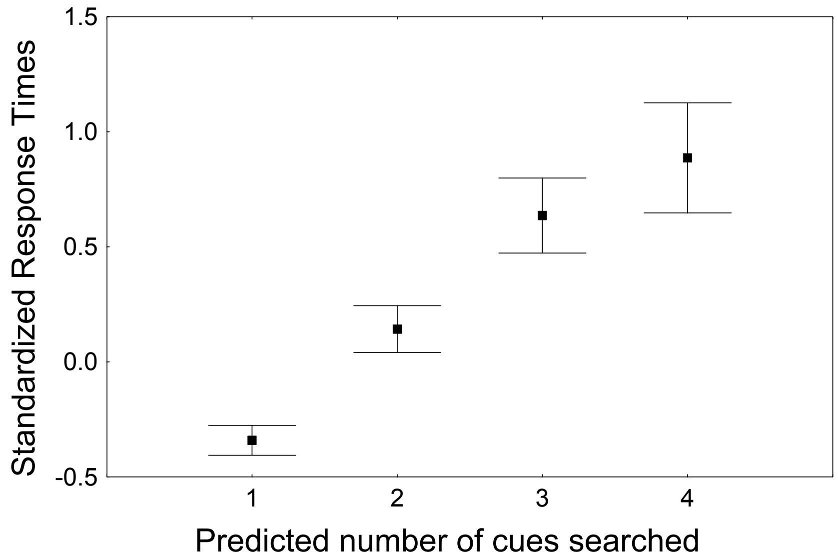

Figure 7. Means of the individually z-transformed response times broken down into groups, as predicted by fitting TTB to each participant. Error bars are 95% confidence intervals for the means. The data is from ten participants in the condition with a TTB-friendly cue structure and the Text presentation format. TTB predicts a linear pattern and the results are very close to that within the error of measurement.

SIMPLEX predicts that decisions are faster when one or two criterions can be estimated by direct retrieval from memory. That is, if the exemplars on the computer screen match exemplars in memory the participant should respond faster. A problem with this prediction is that it will only work well when participants have not yet learned all 16 exemplars. Fortunately most participants did not reach optimal performance with the EX-Friendly cue structure, and the problem was rather that many probably did not use any successful strategy at all. Comparing human response time data with the predictions of SIMPLEX showed an effect at the group level for the EX-Friendly/Letter group as expected but also unexpectedly in the TTB-Friendly/Letter group. The paradox that both TTB and SIMPLEX explained variance for the TTB-Friendly/Letter group was resolved by splitting the group in two. It was then found that the half with the best performers used TTB and the other half was best explained with SIMPLEX. The Letter presentation format may have induced exemplar use unexpectedly with the TTB-Friendly cue structure and then further learning did not occur for unknown reasons.

4.3.2 Experiment 2: single object training and accuracy

Experiment 1 showed that participants were not able to adapt to the environment when exemplars from the EX-Friendly cue structure was presented with the Text format. In Study 1 and 2 the task was assumed to involve covert mental operation on all cues involved and one explanation to why exemplars was not used is that they need to be prelearned. Experiment 2 was designed to give the participants a chance to learn as much as possible about the environment prior to the pair-comparison task.

A single object estimation task was added as a training phase before the pair-comparison task for the experiment groups. The single object estimation task was simply to guess the number of sold vacuum cleaners for a city. A matrix of numbers in boxes were provided beneath the city and the participant simply had to click a box with a number and then feedback was given. The box turned green if it was the correct number. Otherwise the box turned red and the box with the correct number turned green. All 16 exemplars were shown 13 times for a total of 208 trials. The controls had no training before the pair-comparison task and were also used to replicate the Text conditions of Experiment 1. There were four groups with 10 participants each: TTB-Friendly/No Training, TTB-Friendly/Training, EX-Friendly/No Training and EX-Friendly/Training. All groups used the Text presentation format.

The predictions were that single object training would increase accuracy with the EX-Friendly cue structure and that accuracy would decrease with training with the TTB-Friendly cue structure. The latter hypothesis relied on the assumption that exemplars would be used during single object training and the use of TTB in the pair-comparison task should then be blocked or delayed.

Accuracy

The TTB-Friendly/Training group was better in Block 1 compared to the TTB-Friendly/No training group. But there was almost no improvement from Block 1 to Block 2. Thus training had an effect but it did not make the difficult pair-comparison task simpler.

The prediction that single object training would be detrimental to accuracy with the TTB-Friendly cue structure had no support. The accuracy was even closer to optimal behavior and TTB explained even more of the variance for the response times. Note however that the increases in accuracy and explained variance were not significant since the means were both very close to the theoretical maximum.

Model fits of TTB and SIMPLEX

TTB was again found to explain response time data. Accuracy was overall for the EX-Friendly cue structure compared to Experiment 1 higher, but SIMPLEX was not able to explain any significant amount of variance for the response times.

4.3.3 Conclusions

TTB was used easily in the Text condition, whereas exemplar based decision making required the Letter presentation format. This result does not support that humans are boundedly rational in a strong sense, such that we quickly will adapt to any decision environment no matter how the information is presented. We are perhaps not equipped with numerous specialized psychological mechanisms, but rather with a limited number of general mechanisms that strike a balance between accuracy and complexity. The results tentatively support that we have at least two different modes of decision making that work well under certain circumstances, but that there are situations where both modes of decision making have trouble.

This study used a very simple exemplar model, SIMPLEX for many reasons as detailed in Study III. The main reason was that the simple model would predict in most cases predict the same decisions as PROBEX for the cue structures used in the experiments. As a consequence all empirical conclusions about exemplar based reasoning hold are valid for the assumption of exemplar representation. Similarity based reasoning is not tested at all in these experiments, except that the response times predicted by SIMPLEX are based on that direct retrieval of an identical exemplar is assumed to be faster than retrieving a similar exemplar.

These results are also the first where most participants used TTB with good accuracy, without any incitement such as time pressure or a cost attached to investigating each feature.

5 DISCUSSION

This summary has argued that decision theories of bounded rationality should incorporate basic theories of cognition and perception, and that it is crucial to consider the information structures of the environments. The idea of bounded rationality is old, but has had little impact on theories of decision making.

This thesis presented PROBEX which is a boundedly rational process model of inference and probability judgment. This model was based upon very successful exemplar based models for perceptual categorization (Medin & Schaffer, 1978; Nosofsky, 1986). It was shown in computer simulations that PROBEX is very competitive compared to other decision strategies (Study I, II) with almost no knowledge as well as with unlimited knowledge. It was fitted to human data from a general knowledge task and was able to fit the point estimates, decision and probability assessments with parameters that indicated that the effort was both ecologically rational and psychologically plausible (Study I). It was also shown that it theoretically is more important to make good decisions early in learning than later from an evolutionary point of view (Study II).

A simplified version of PROBEX, SIMPLEX, was used along with Take The Best (Gigerenzer & Goldstein, 1996) to design two cue structures used to evaluate whether decision makers would adapt to these cue structures. It was found that participants easily adopted TTB in a noncompensatory environment if the cues were presented with text labels and that SIMPLEX could be used in a difficult nonlinear environment when cues were easily processed as short strings of letters (Study III).

These results support the idea of ecological rationality, that is, efficient and accurate human decision making, as long as the information is structured and presented in a way suitable for some decision mechanism.

This thesis covered tasks where two objects are compared and one is selected on the basis of an unknown feature, the criterion. Decision making is much more complex and can be applied to, for example, foreign policy (Redd, 2002) as well as to the selection of a new toothbrush at the supermarket, but hopefully knowledge of basic decision processes should generalize as building blocks to theories of more complex decision situations.

5.1 Process Models of Decision Making

Process models are important in the categorization literature. The Generalized Context Model (Nosofsky, 1986), neural networks (Kruschke, 1992), decision bound theory (Ashby, Alfonso-Reese, Turken & Waldron, 1998), prototype models (Smith & Minda, 1998) to name a few all contribute to a healthy exchange of scientific ideas.

However, the study of processes in decision making has not stressed computational modeling. Payne et al. (1993) did examine possible tradeoffs between accuracy and effort for several decision strategies with computer simulations, but they have not used computational modeling to derive and test hypothesis as is customary in the categorization literature. The Adaptive Toolbox project (Gigerenzer & Todd, 1999) recently launched an initiative to investigate models of bounded rationality and it is only through empirical studies using computational modeling that these well defined models of decision making can be corroborated. PROBEX could be seen as a tool that also belongs to the adaptive toolbox and there are now a small set of psychologically plausible theories of decision making ripe for empirical investigation.

Study I and II were mostly theoretical, although PROBEX was shown to be consistent with human behavior by fitting it to decision data. Study III was an attempt to make predictions from TTB and a simplified version of PROBEX and test these predictions empirically. Model fits that predicited response times was used to investigate the data to find traces of the processing required for exemplar based and one reason decision making.

The first premise of this discussion is that process modeling is relevant to decision making as it is evidently relevant to categorization. But is it obvious that this is the case? Categorization processes could be fundamental to the human mind, whereas decision making is contingent (Payne et al., 1993) on situations to the extent that decision strategies are not predictable processes. If important brain systems are involved in decision making then it should be possible to model decision making when carefully chosen cue structures are used. The second premise is that the idea of bounded rationality can increase the power of theory development and empirical research by the study of cue structures.

5.1.1 Conclusions for exemplar-based decision making

There is a *strict match* when cues of a probe are identical to the cues of a mental representation. Hahn and Chater (1998) distinguish between partial and strict matching of cues to separate similarity-based processes from memory processes. PROBEX is a prime example of a similarity-based process, whereas SIMPLEX is more inclined towards a memory process where similarity based retrieval is used only when direct retrieval is impossible.

Study I gave support for similarity-based reasoning by fitting PROBEX to human decision data, but, of course, it cannot be ruled out that some other model would explain this data better. Only the names of the cities were given to participants. The support for similarity-based reasoning comes from the assumption that the cue structure used in the model was a good approximation of the cue structure used by the participants.

Study III supported exemplar-based decision making but cannot be seen as a test of similarity-based decision making, but it did support the use of exemplars when the alternatives were presented as easily encoded strings of letters.

The empirical support for exemplar based decision making has only recently begun to assemble (Dougherty et al., 1999; Juslin et al., in press; Juslin, Nilsson & Olsson, 2001; Juslin, Olsson & Olsson, 2003; Sieck & Yates, 2001), but it is clear that this fresh approach to decision making is making progress.

There is an omission in PROBEX and a problem for lazy algorithms in general that should be addressed in conjunction with the results of study III. It is common in most exemplar based models to have attention parameters for each feature/dimension. This was not included in the original model because it was supposed that the model would foremost apply to tests of general knowledge without prior training (most people have never compared the size of German cities). Lazy algorithms suffer from the "dimensionality curse" which is a problem when large data sets are used and most dimensions are not informative at all to the task at hand (Aha, 1997; Hahn & Chater, 1998). There is no computationally simple way to know which dimensions should be attended to or not. A human might use common sense to do so, but common sense seems not to be a natural part of an automatic process such as exemplar retrieval. In study III most participants failed to learn the EX-Friendly cue structure with the Text presentation format in Experiment 1, and the training in Experiment 2 did not help much. A possible explanation is that we avoid encoding exemplars because of the dimensionality curse, until we know for sure which features are worth attending to. This seems to be consistent with the representational shift from simple rules to exemplar use later in training that Johansen and Palmeri (2002) found in a number of categorization experiments. An important issue in the future is to understand when and how exemplars are learned.

5.1.2 Conclusions for one reason decision making

An *abstraction* represents a class of objects rather than an instance of a class. A simple example is when the abstraction only uses a few cues and ignores other cues. Hahn and Chater (1998) define rule based decision making as processes where cues are strictly matched to abstractions, and TTB fits this definition when it makes the decision. The abstraction of one cue is strictly matched. But TTB is also a serial search process and compared to ordinary rules this implies strong predictions about response times. Study III found clear evidence of such response time patterns when the processing of cues required effort. In the condition where cues were presented as short strings of letters rather than text labels, many participants failed to adapt to the task.

The obvious objection to these otherwise strong results is that the cue structure itself reinforced the use of TTB. It is possible that using a probabilistic task rather than a deterministic task would be much harder.

There might also be other models that can predict these response times as well as TTB does. To my knowledge there is no such model. For example other categorization models as EBRW (Nosofsky & Palmeri, 1997) and EGCM-RT (Lamberts, 2000) categorize each object in a pair comparison independently, while the response times patterns of TTB comes from searching both objects simultanously making an intricate pattern that cannot be emulated by summing up two independent response times.

5.1.3 Adapted to adapt

Evolutionary psychology (Cosmides & Tooby, 1994) states that cognition is highly modular and that the modules are adapted to specific tasks. An alternative view is that the human mind evolved to adapt to environments during the course of a lifetime. The implication is that evolution provided the human mind with either hardwired decision strategies or low level mechanisms that adapts to decision tasks during a lifetime.

The experiments in Study III showed huge variations between subjects. This is more consistent with the idea that decision makers learn strategies. But is there then any use for process models of decision making? Even if there are no evolutionary evolved modules or strategies for decision making, process models can still be used to understand how different cue structures may be handled.

Decision making may also gain benefits from process models by directly modeling learning. SUSTAIN (Love et al., in press) is a neural network model of categorization that is neither rule-based nor similarity-based within the process taxonomy of Hahn and Chater (1998). It is a frugal learning model to borrow a term from the adaptive toolbox, since it starts out with as few clusters of neurons as possible and then add new clusters when it is unable to solve problems correctly. SUSTAIN has been used to correctly predict the difficulty of common cue structures used in the categorization literature. A process model applied to learning opens up new possibilities. Clapper and Bower (2002) showed that the sequencing of training instances in unsupervised categorization had a strong effect on the learning rate, and with SUSTAIN it could be possible to make testable predictions from the sequencing of training exemplars. The sequences in Study III were randomized for each participant, which do not reflect reality. If the cues of the environment have structure, then the sequence of real world trials are likely to be ordered rather than random.

5.2 Final Remarks

Bounded rationality is easy to misinterpret as irrationality, but given the complexities of the world a simple strategy may be perfectly rational because it is not possible to do better given the constraints of our mental capacities. The implication is not that humans always are rational (human rationality has been hotly debated many times (e.g. Cohen, 1981; Jungermann, 1983; Stanovich & West, 2000), but simply that in some situations we are able to make good decisions. This dissertation is an example of how formal modeling perhaps is a necessary tool in order to make real progress in empirical psychology. In the case of decision making the results here hint that time is ripe to unify decision making theory and recent advances in cognitive psychology. Yet, the ideas presented here are to a great extent still only ideas so far. The perhaps most important idea is to incorporate the cue structures of the environment into experimental design, but study III only used two highly artificial cue structures that are unlikely to be found outside the laboratory. Yet, this technique covered a much wider range of possible cue structures than is typical in similar experiments.

REFERENCES

Aha, D. W. (1997). Lazy learning. *Artificial Intelligence Review, 11,* 7-10.

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

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science, 2,* 396-408.

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A formal neuropsychological theory of multiple systems in category learning. *Psychological Review, 105,* 442-481.

Augier, M. (2001). Sublime Simon: The consistent vision of economic psychology's Nobel laureate. *Journal of Economic Psychology, 22,* 307-334.

Beach, L. R., & Mitchell, T. R. (1987). Image theory: Principles, goals, and plans in decision making. *Acta Psychologica, 66,* 201-220.

Bröder. A. (2000). Assessing the empirical validity of the "Take-The-Best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 1332-1346.

Chater, N. (2000). How smart can simple heuristics be? *Behavioral and Brain Sciences, 23*, 745-746.

Clapper, J. P., & Bower, G. H. (2002). Adaptive categorization in unsupervised learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28,* 908-923.

Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences, 4,* 317-370.

Cosmides, L., & Tooby, J. (1994). Beyond intuition and instinct blindness: toward an evolutionary rigorous cognitive science. *Cognition, 50,* 41-77.

Daelemans, W. (1995). Memory-based lexical acquisition and processing. In P. Steffens (Ed.), *Machine Translation and the Lexicon*, (Springer Lecture Notes in Artificial Intelligence No. 898, pp. 85-98). Berlin: Springer.

Damasio, A. R. (1998). Emotion in the perspective of an integrated nervous system. *Brain Research Reviews, 26,* 83-86.

Doherty, M. E., & Brehmer, B. The paramorphic representation of clinical judgment: A thirty-year retrospective. In W. M Goldstein & R. M. Hogarth (Eds.), *Research on judgment and decision making. Currents, connections and controversies* (pp. 537-551). Cambridge: Cambridge University Press.

Dolan, R. J. (2002). Emotion, cognition, and behavior. *Science, 298,* 1191-1194.

Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory processes model for judgments of likelihood. *Psychological Review, 106,* 180-209.

Einhorn, H. J., Kleinmuntz, D. N., & Kleinmuntz, B. (1979). Linear regression and process-tracing models of judgment. *Psychological Review, 86,* 465-485.

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

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review, 103*, 650-669.
Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 3-34). Oxford: Oxford University Press.
Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: the recognition heuristic. *Psychological Review, 109,* 75-90.
Goldstein, W. M., & Hogarth, R. M. (1997). Judgment and decision research: some historical context. In W. M Goldstein & R. M. Hogarth (Eds.), *Research on judgment and decision making. Currents, connections and controversies* (pp. 537-551). Cambridge: Cambridge University Press.
Hahn, U., & Chater, N. (1998). Similarity and rules: distinct? exhaustive? empirically exhaustive? *Cognition, 65,* 197-230.
Hammond, K. R. (2000). Coherence and correspondence theories in judgment and decision making. In T. Connolly, H. R. Arkes, and K. R. Hammond (Eds.), *Human Judgment and Social Policy.* New York: Oxford University Press.
Hanoch, Y. (2002). "Neither an angel nor an ant": Emotion as an aid to bounded rationality. *Journal of Economic Psychology, 23,* 1-25.
Harte, J. M., Westenberg, M. R. M., & van Someren, M. (1994). Process models of decision making. *Acta Psychologica, 87*, 95-120.
Hastie, R. (2001). Problems for judgment and decision making. *Annual Review of Psychology, 52,* 653-683.
Hintzman, D. L. (1988). Judgments of frequency and recognition memory in a multiple trace memory model. *Psychological Review, 95,* 528-551.
Huber, O. (1994). Simulation as a tool to improve process models. *Acta Psychologica, 87,* 121-134.
Johansen, M. K., & Palmeri, T. J. (2002). Are there representational shifts during category learning? *Cognitive Psychology, 45,* 482-553.
Jungermann, H. (1983). The two camps on rationality. In Scholz, R. W. (Ed.), *Decision Making under Uncertainty* (pp. 63-86). Amsterdam: North-Holland.
Juslin, P., Jones, S., Olsson, H., & Winman, A. (in press). Cue abstraction and exemplar memory in categorization: Evidence for multiple representation levels. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*
Juslin, P., Nilsson, H., & Olsson, H. (2001). Where do probability judgments come from? Evidence for similarity–graded probability. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society.* (pp. 471-476). Mahwah, NJ: Erlbaum.
Juslin, P., Olsson, H., & Björkman, M. (1997). Brunswikian- and Thurstonian-origins of bias in probability assessment: On the interpretation of stochastic components of judgment. *Journal of Behavioral Decision Making, 10*, 189-209.
Juslin, P., Olsson, H., & Olsson, A-C. (2003). Exemplar effects in multiple-cue judgment. *Journal of Experimental Psychology: General, 132,* 133-156.
Juslin, P., & Persson, M. (2002). PROBabilities from EXemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge. *Cognitive Science, 26*, 563-607.
Juslin, P., Wennerholm, P., & Winman, A. (1999). Format-dependence in subjective probability calibration. *Journal of Experimental Psychology: Learning, Memory and Cognition, 25,* 1038-1052.

Kahneman, D. (2000). A psychological point of view: violations of rational rules as a diagnostic of mental processes. *Behavioral and Brain Sciences, 23,* 681-683.

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

Kasif, S., Salzberg, S., Waltz, D., Rachlin, J., & Aha, D. W. (1998). A probabilistic framework for memory-based reasoning. *Artificial Intelligence, 104,* 287-311.

Klein, G. A., Calderwood, R., & Macgregor, D. (1989). Critical decision method for eliciting knowledge. *IEEE Transactions on Systems, Man, and Cybernetics, 19,* 462-472.

Klein, N. M. (1983). Utility and decision strategies: a second look on the rational decision maker. *Organizational Behavior and Human Performance, 31,* 1-25.

Klemm, K., Bornholdt, S., & Schuster, H. G. (2000). Beyond Hebb: exclusive-OR and biological learning. *Physical Review Letters. 84,* 3013-3016.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning, *Psychological Review, 99,* 22-44.

Lamberts, K. (2000). Information-accumulation theory of speeded categorization, *Psychological Review, 107*, 227-260.

LeDoux, J. E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience, 23,* 155-184.

Lipshitz, R., Klein, G., Orasanu, J., & Salas, E. (2001). Taking stock of naturalistic decision making. *Journal of Behavioral Decision Making, 14,* 331-352.

Logan, G. D. (1988). Towards an instance theory of automatization, *Psychological Review, 95,* 492-527.

Logan, G. D. (2002). An instance theory of attention and memory. *Psychological Review, 109,* 376-400.

Lopes, L. (1991). The rhetoric of irrationality. *Theory & Psychology, 1,* 65-82.

Lopes, L. (1992). Three misleading assumptions in the customary rhetoric of the bias literature. *Theory & Psychology, 2,* 231-236.

Love, B. C., Medin, D. L., & Gureckis, T. M. (in press). SUSTAIN: a network model of human category learning. *Psychological Review.*

Martignon, L., & Laskey, K. B. (1999). Bayesian benchmarks for fast and frugal heuristics. In G. Gigerenzer, P. M. Todd, and the ABC Research Group. *Simple heuristics that make us smart* (pp.169-188). New York: Oxford University Press.

Medin, D. L., & Schaffer, M. M. (1978). Context model of classification learning. *Psychological Review, 85,* 207-238.

Mellers, B. A. (2000). Choice and the relative pleasure of consequences. *Psychological Bulletin, 126,* 910-924.

Mellers, B., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of Experimental Psychology: General, 128,* 332-345.

Montgomery, H. (1989). From cognition to action: The search for dominance in decision making. In H. Montgomery & R. O. Svenson (Eds.), *Process and structure in human decision making.* Chichester: Wiley.

Montgomery, H., & Svenson, O. (1976). On decision rules and information processing strategies for choices among multiattribute alternatives. *Scandinavian Journal of Psychology, 17,* 283-291.

Newell, A. (1992). Précis of Unified theories of cognition. *Behavioral and Brain Sciences, 15,* 425-492.

Newell, B. R., & Shanks D. R. (2003). Take the best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory and Cognition, 29*, 53-65.
Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical test of a fast-and-frugal heuristic: Not everyone "takes-the-best". *Organizational Behavior and Human Decision Processes, 91,* 82-96.
Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39-57.
Nosofsky, R. M., Kruschke, J. K., & McKinley, S. C. (1992). Combining exemplar-based category representation and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 211-233.
Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review, 104,* 266-300.
Palmeri, T. J., & Flanery, M. A. (2002). Memory systems and perceptual categorization. *Psychology of Learning and Motivation: Advances in Research and Theory,41,* 141-189.
Payne, J. W., & Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge: Cambridge University Press.
Persson, M. (2003). Decision strategies as adaptations to cue structures. Manuscript in preparation.
Persson, M., & Juslin, P. (2000). Fast and frugal use of cue direction in states of limited knowledge. In L. R. Gleitman & A. K. Joshi (Eds.), *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society* (pp. 841-846). Mahwah, NJ: Erlbaum.
Pitt, M. A., Myung, I. J., & Zhang, S. B. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review, 109,* 472-491.
Rachlin, J., Kasif, S., Salzberg, S., & Aha, D. W. (1994). Towards a better understanding of memory-based reasoning systems. In H. Hirsch & W. Cohen (Eds.), *Machine Learning: Proceedings of the Eleventh International Conference*, (pp. 242-250). San Fransisco: Morgan Kaufmann.
Redd, B. R. (2002). The influence of advisers on foreign policy decision making – an experimental study. *Journal of Conflict Resolution, 46,* 335-364.
Roberts, S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review, 107,* 358-367.
Sieck, W. R., & Yates, J. F. (2001). Overconfidence effects in category learning: A comparison of connectionist and exemplar memory models. *Journal of Experimental Psychology: Learning, Memory and Cognition, 27,* 1003-1021.
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.
Simon, H. A. (1992). Rational decision-making in business organizations: Nobel memorial lecture, 8 December 1978. In A. Lindbeck (Ed.), *Nobel Lectures, Economic Sciences 1969-1980* (pp. 343-371). Singapore: World Scientific Publishing.
Sloman, S. A., & Rips, L. J. (1998). Similarity as an explanatory construct. *Cognition, 65,* 87-101.
Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review, 99,* 3-21.

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1411-1436.

Stanovich, K. E., & West, R. F. (2000) Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences, 23,* 645-726.

Storms, G., De Boeck, P., & Ruts, W. (2000). Prototype and exemplar-based information in natural language categories. *Journal of Memory and Language, 42,* 51-73.

Svenson, O. (1992). Differentiation and Consolidation Theory of human decision making: A frame of reference for the study of pre- and post-decision processes. *Acta Psychologica, 80,* 143-168.

Svenson, O. (1996). Decision making and the search for fundamental psychological regularities: What can be learned from a process perspective? *Organizational Behavior and Human Decision Processes, 65,* 252-267.

Todd, P. M., & Gigerenzer, G. (1999). What we have learned (so far). In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 3-34). Oxford: Oxford University Press.

Tversky, A. (1977). Features of similarity. *Psychological Review, 84,* 327-352.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science, 185,* 1124-1131.

Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review, 90,* 293-315.

ACKNOWLEDGEMENTS

The undisputed champion of scientific writing is my supervisor Peter Juslin. He knows what it takes to make a difficult subject readable. His creativity is also unending, no matter how busy he is, he always has something new to discuss. He has also been patient with me and supported me when it mattered most. The second next important person for the content of this thesis is Henrik Olsson, the 'human encyclopedia' who commented my texts on all stages and forced me to think harder on numerous occasions. A third person is also very important. Without the support and advice from Berit Hagekull this would not have been written. She has been my mental discipline coach, or in a more psychodynamic vocabulary: my externalized superego. When it comes to close support, Gustaf Gredebäck is the man who has been spinning my chair whenever I fall asleep on duty. Professor em. Mats Björkman was a great inspiration for me, but is sadly no longer with us. The project colleagues Anders Winman and Pia Wennerholm have been helpful and friendly. Sverker Runeson has been an inspiration for me both as a researcher and teacher. The secretary Mildred Larsson has made practical matters easy for me. Mats Olsson and Claes von Hofsten commented earlier writings. Meeting the researchers in the ABC-Group from Berlin has induced some of the most creative periods I have had as a researcher. Hungry men with beards have been company for lunch. All the social gatherings with fellow graduate students during, for example, informal late Friday seminars were important for my wellbeing, as well as all other employees who I normally have met during daytime. The heads of the department were crucial: Outi Lundén for arranging extra funding from the university early on and Erik Börjesson for absentmindedly forgetting to kick me out when the funding ended. Special thanks to Ben Newell who found 52 grammatical errors in the thesis the last day. Lars-Erik Larsson is the man who keeps the department from falling apart. There has been a life (too much?) outside of the department too. My family has been very patient and still wonder when I will grow up. The society of Gôssegôbbera has made my beer selection strategies much easier, and the other (more attractive) half of "the choir of Värmlands nation " has stimulated some nonacademic writings. The game of Go is both the reason why I became interested in psychology and the reason why this took so long time.

Acta Universitatis Upsaliensis

Comprehensive Summaries of Uppsala Dissertations
from the Faculty of Social Sciences
Editor: The Dean of the Faculty of Social Sciences

A doctoral dissertation from the Faculty of Social Sciences, Uppsala University, is either a monograph or, as in this case, a summary of a number of papers. A few copies of the complete dissertation are kept at major Swedish research libraries, while the summary alone is distributed internationally through the series *Comprehensive Summaries of Uppsala Dissertations from the Faculty of Social Sciences*. (Prior to July 1985, the series was published under the title "Abstracts of Uppsala Dissertations from the Faculty of Social Sciences".)

Distributor:
Uppsala University Library,
Box 510, SE-751 20 Uppsala, Sweden
www.uu.se, acta@ub.uu.se

ISSN 0282-7492
ISBN 91-554-5697-9