PREFACE

Everyone writing a dissertation is faced with the problem of what level of detail to choose when explaining to family, friends, and acquaintances what the dissertation is about. I usually opted for a very general answer, saying that I am studying how people make decisions under various circumstances. To my surprise, many of my conversation partners got very excited hearing that answer and told me about a troubling decision problem they had faced in the past, such as whether or not to buy a house, what profession to learn, or which of two very attractive job offers to accept. As these examples show, mostly high-stakes preferential choices were mentioned. Any reader interested in these kinds of decisions will be just as disappointed as my vis-à-vis in these situations: This is what my dissertation will *not* be about. Instead, it will highlight the small decisions we make every day, sometimes without even noticing, and it will present a possible reason why when making these decisions we mostly do not experience the psychological pains, efforts, and doubts we remember of the important crossroads in our life, that is, the use of fast and computationally simple decision heuristics.

Moreover, the focus will be on inferences rather than preferences. For inferences, in contrast to preferences, an objective outside criterion exists on which to evaluate whether a decision has been correct. Thus, while there is no accounting for taste, the correctness of an inference can be determined objectively. As much of my work is concerned with the *accuracy* of decision strategies, this feature makes inferences a very useful task to study. But what is the relevancy of studying inferences when the decisions that people seem to care about most are preferential choice problems? Apparently caring more about preferential choice problems than about inferences does not mean that the former are more common or more important. Many everyday decisions involve inferences, such as inferring at which store we will most likely find a certain product or by which means of transport we will get to the airport most quickly. We might sometimes not be aware of these decisions because most are done quickly and without much mental effort, maybe exactly because mental shortcuts – heuristics – are at work. Such shortcuts can also be applied to preferential choices, and I will refer to work on this topic whenever a connection can be made. Yet a feature more commonly associated with preferential choices as opposed to inference problems is the low or even negative correlations between the attributes of the alternatives to choose from, such as the often negative correlation between price and quality of a product. This characteristic can make choices difficult, which might be the reason for their greater salience to us when we think about decision making. The effect of correlations between cues, that is, the degree of information redundancy, on the decision process will be explored in detail for the domain of inferences in the second chapter of my dissertation. But let me give a more comprehensive overview of my work summarized in this dissertation.

In general, my dissertation argues for an adaptive view of decision making. This means that decision makers respond adaptively to characteristics of the decision environment – such as whether the decision has to be made under time pressure – and apply decision strategies that work well under the given circumstances. Specifically, my dissertation focuses on paired comparisons: One has to infer which of two alternatives, described on several cues, has a higher criterion value, such as which of two cities has more inhabitants, which of two apartments is more expensive, or at which of two potential drilling sites more oil is to be found. The task requires that a probabilistic inference is made, as the cues to rely on are positively, but not perfectly correlated with the criterion. The Adaptive Behavior and Cognition (ABC) research group of the Max Planck Institute for Human Development has formulated several simple heuristics for paired comparisons (e.g., Gigerenzer, Todd, & the ABC Research Group, 1999). These heuristics were tested in computer simulations on a variety of real-world datasets and often showed astonishing performance. Relying on only a fraction of the available information, some of these heuristics were able to achieve similar, and under certain circumstances even higher accuracy compared to more complex models, such as multiple regression (Czerlinski, Gigerenzer, & Goldstein, 1999). From an adaptive point of view, the combination of simplicity and accuracy makes these heuristics plausible candidates for human decision processes. Yet in contrast to the impressive demonstrations of the performance of simple heuristics in computer simulations, empirical evidence for their use was initially scarce. In the meantime, the body of evidence has grown. In the first chapter of this dissertation I will explore the conditions under which certain heuristics work well and will review the available empirical evidence for their use. In the second and third chapters, I will report my own simulations and experimental studies.

The study of decision heuristics raises the questions of under which environmental conditions their use is adaptive, and to which of these conditions people adaptively respond by using the appropriate heuristics. To answer these questions, in the first chapter I shift the focus from heuristics to their building blocks. As in the periodic table of elements, there are a large number of heuristics composed of a few building blocks. Three building blocks will be studied in detail: psychological principles of information search, stopping search, and deciding. An ecological analysis will show which building blocks fit which environmental structures. This analysis will, under the general assumption that human decision making is adaptive, allow precise predictions about how heuristics are "selected" by the environment or the experimental task. To test these predictions, empirical evidence, so far available for only some of the derived hypotheses, will be reviewed. It will further be shown that decision rules depend on the kind of stopping rules used, but not on search rules. These constraints pose limits to the combination of building blocks, and therefore, to the number and kinds of possible heuristics. In sum, the first chapter will thus provide a theoretical framework, structuring how the interplay between heuristics and the decision environment can be empirically studied. How such an analysis is conducted precisely will be demonstrated in the

second chapter for one of the predictions of the ecological analysis put forward in the first chapter.

More specifically, in the second chapter I will present three studies that examine the effect of information redundancy (measured via the correlations between cues) on the performance and the use of different decision strategies for probabilistic inferences. A simulation study will compare different decision strategies under high versus low information redundancy in terms of both their frugality and their accuracy. From the results of the computer simulation predictions will be derived about which strategies people use when they encounter high- versus low-redundancy decision environments. Consequently, two experiments will address the question of whether people adaptively select strategies in response to both information costs and information redundancy. In the first experiment, people will have the opportunity to learn about the differential accuracy of strategies through outcome feedback, whereas in the second experiment this opportunity will be removed, thus exploring whether a more intuitive form of adaptivity to environmental circumstances also exists. An interesting side question addresses individual differences as a second source of variance, besides differences in the structure of the decision environment. More precisely, the association of discomfort with ambiguity and the use of simple non-compensatory strategies as a means to avoid ambiguity in low-redundancy environments is investigated.

The third chapter takes one step back. So far, the set-up of a heuristic's building blocks and what it takes to be able to apply them has not been given much attention. Instead, in simulations a certain building block such as the search rule has mostly been assumed as a given, and in experiments participants were usually told the order in which to look up information. Indeed, it has been criticized that simple decision heuristics such as Take The Best owe much of their simplicity and success to the considerable computations necessary for setting up the cue search order before the heuristic can be used. In the third chapter I will address this criticism. First, it will be shown that there are typically many cue orders that will achieve good performance in a given problem domain. Second, simple learning rules will be suggested that step-by-step reorder cues in the online process of decision making based on outcome feedback. In simulations, it will be tested whether these rules are able to converge on one of the many useful cue orders through exposure to just a small number of decisions. Finally, an experimental study will explore to what extent the proposed simple learning rules can account for the way participants adapt their search through cues in a decision making task when they do not have prior knowledge about the differential usefulness of these cues.

A final section discusses what my dissertation can add to a more comprehensive theory of an adaptive toolbox filled with simple heuristics, while also critically pointing out which questions will need further exploration. Overall, the reported studies will on the one hand provide evidence that human decision making indeed is adaptive. On the other hand, it will also be shown that there are limits to adaptivity and that some of the assumptions in the fast and frugal heuristics program might need to be revised, to contribute to a more realistic picture of human decision making.