Summary and Concluding Remarks

My goal in this dissertation is to argue that people are equipped with multiple decision strategies that they apply adaptively in response to various circumstances. The focus is on building blocks of heuristics because these allow for high flexibility while, at the same time, preserving parsimony: There are a large number of heuristics composed of only a few building blocks. Focusing on paired comparisons, I derived precise predictions about when which building blocks will be selected, tested some of the predictions empirically, and finally also put to a test basic assumptions on which simple heuristics, such as Take The Best, rest. Before discussing some of the implications of the work presented in this dissertation, I will briefly summarize the main findings.

In the first chapter, I analyzed key features of decision environments in terms of which heuristic building blocks of search, stopping search, and deciding they support and favor. From this ecological analysis, I derived a list of hypotheses about when which building block should be used, under the general assumption that human decision making is adaptive. Some of these hypotheses have already gained support in the experiments that were reviewed in the first chapter: One-reason stopping and deciding predicts people's decision processes and outcomes well when information about cues is costly, when decisions have to be made under time pressure, and when information has to be retrieved from memory. When information is relatively cheap, people more often continue to search beyond a first discriminating cue and integrate the gathered information in a compensatory fashion. The exact stopping rules that are used in these situations need to be explored further, but preliminary evidence speaks for simple tallying rules (e.g., to stop when two reasons have been encountered that support one alternative) instead of exhaustive search and complicated information integration models. The first chapter also pointed out problems with previous experiments on the use of simple inference heuristics. Hence, several hypotheses, especially those concerning search rules, still have to be tested in a more stringent way. There is evidence that people, under many circumstances, follow a certain search order when looking up cues on which to base their decision. The exact ordering criterion they use, however, and whether it changes according to environmental characteristics remains unclear, mostly due to confounds in the respective experiments. Also, the question of how people learn cue orders, especially an order by validity, has not been answered in previous studies because participants were mostly told in which order to best look up information. For this reason, I picked up the issue of search rules and how they are constructed in the last chapter of my dissertation.

But prior to this, one of the hypotheses derived in the first chapter was explored in detail in the second chapter. The prediction was that a high degree of information redundancy (as measured by the intercorrelation between cues) favors one-reason stopping and decision making. Indeed, in a computer simulation, I could show that different strategies achieve highly similar accuracy in redundant environments, which makes the frugal one-reason decision making heuristic Take The Best very attractive in terms of cost-benefit ratio. Only when both degree of redundancy and cue validity dispersion were low did compensatory strategies clearly outperform Take The Best in terms of accuracy. Simple compensatory strategies, however, did not fall far behind more complex ones. In two experiments, I investigated whether people adaptively respond to the degree of redundancy of the decision environment. The results indicate that people indeed adaptively select strategies in response to both information costs and information redundancy. In particular, process measures support the notion that Take The Best is used predominantly in a highly redundant decision environment in which cues correlate positively with each other, whereas in a low-redundancy environment, compensatory, yet still frugal, strategies are used. Interestingly, adaptive strategy selection could still be demonstrated when the possibility to learn from outcome feedback was removed, suggesting that even without being able to find out about the differential performance of strategies, certain environmental features function as triggers for the use of certain strategies.

Furthermore, my results emphasize the role of individual differences in decision-making processes. Undoubtedly, environmental characteristics can explain much of the variance observed in decision processes. But nevertheless, people differ in the strategies they select in response to certain environments. In the case of low information redundancy, at least when no outcome feedback was provided, participants seemed to systematically differ in their willingness to face conflicting information. Those who reported high discomfort with ambiguity were more likely to apply one-reason stopping and decision making (which avoids encountering conflicting information) when they face a decision environment with slightly negative correlations between cues.

In the last chapter, I addressed the question of how a particular building block, the search rule, can be constructed in the first place. Previous simulation studies were based on precomputed cue orders by, for example, ecological validity (as in Take The Best). Similarly, in experimental studies participants were mostly told in which order to look up cues. Yet, this simplifying assumption is questioned on grounds of psychological plausibility: Calculating ecological validities is expensive, both in terms of information storage as well as computational demands. Generally agreeing to the necessity of taking into account the set-up costs of a heuristic in addition to its application costs when considering its overall simplicity, I proposed simple cue order learning rules and tested them in simulations to address this justified criticism. The cue orders resulting from the application of these learning rules approach Take The Best's accuracy, while being much more frugal. Thus, in principle, simple learning rules enable a one-reason decision heuristic to, simultaneously, become psychologically more plausible and perform only slightly worse than if it had full knowledge of cue validities from the very beginning. The results of a subsequent experiment support the notion that some of the proposed learning rules are psychologically valid descriptions of the participants' cue ordering process. The rules that reorder cues based on continuously updated simple unconditional tallies of, for instance, the number of correct decisions made by each cue so far fit participants' behavior quite well. In harsh contrast, a rule that reorders cues based on the ratio of the number of correct decisions to the number of discriminations per cue (i.e., its current validity) fits hardly better than a model that assumes that a random order is used. Thus, the results challenge the assumption that people can easily construct a cue order based on validity, on which the most prominent simple heuristic proposed by the ABC research group – Take The Best – rests.

What can these results contribute to the previous work of the ABC research group on simple decision making heuristics, summarized in the metaphor of an *adaptive toolbox* (e.g., Gigerenzer & Todd, 1999)? Some aspects of this work have received much criticism by other researchers in the field of judgment and decision making. While the often repeated criticism that the proposed simple heuristics lack empirical support (Bröder, 2000, 2001, 2003; Bröder & Schiffer, 2003; Oaksford, 2000) could have been abandoned already some time ago in light of growing evidence gathered to a substantial part by the major critic – Arndt Bröder – himself, other issues remained valid. Let me focus on three main points that were raised: (1) Cooper (2000) noted that precise predictions about when which heuristic is applied were lacking, compromising the falsifiability of the adaptive toolbox approach. A similar point was raised by Newell et al. (2003). (2) Another critical issue is that the question of how heuristics are selected from the toolbox remained unaddressed (Feeney, 2000; Luce, 2000; Morton, 2000; Newstead, 2000; Shanks & Lagnado, 2000; Wallin & Gärdenfors, 2000). (3) Finally, the question of how building blocks, especially the search order, can be derived in simple ways (Juslin & Persson, 2002; Wallin & Gärdenfors, 2000) or, more generally, of where the heuristics in the adaptive toolbox come from (Hammerstein, 2000; Luce, 2000) were left without empirical investigation.

Indeed, the book *Simple heuristics that make us smart* (Gigerenzer et al., 1999) might have been relatively silent about these issues. In my dissertation I tried to at least address some of these important questions. In the following section, I will pick up each of the main criticisms summarized above, and show how my work could serve to reply to these criticisms.

First, in order to derive falsifiable predictions, one needs to specify when certain heuristics are applied and when not. Although some approaches in this direction were made (e.g., Martignon & Hoffrage, 1999; Rieskamp & Hoffrage, 1999), maybe the general adaptive point of view on the fit between environments and heuristics should have been stressed even more. This could have possibly avoided misconceptions of the sort exemplified by Bröder (2000, Experiments 1 & 2) who tested Take The Best in conditions for which it is not suited (simultaneously available cues), Newell et al. (2003) who failed to perform a thorough ecological analysis on how well different heuristics would perform in their experimental environments, or Oppenheimer's (2003) work on the recognition heuristic who assumed that it is supposed to be applied regardless of the structure of the environment. But, to take an example from a craftsman's instead of a decision maker's toolbox, the fact that a hammer and nail work well when putting up a picture on a wooden wall does not imply that these are good tools for the same task when the wall is made of concrete. The first chapter worked out for several building blocks under which circumstances they perform well, and where their limits are. The chapter further tried to clarify that environmental characteristics provide information about whether certain building blocks work well without the necessity to compute the optimal search order, stopping point, or decision rule through exhaustive cost-benefit analysis.

Nevertheless, and this brings us to the second point, predictions about how exactly heuristics are selected from the toolbox are needed to advance the theoretical significance of the approach. The claim is that the selection of adequate heuristics from the adaptive toolbox is not just a form of optimization under constraints. But how else is it accomplished? Rieskamp and Otto (2004) propose general mechanisms of reinforcement learning at the level of complete strategies as one possible answer. Yet, in the second chapter, data is presented that suggests that there are also more intuitive forms of adaptivity that might rely on the anticipated performance of decision strategies. Some environmental conditions are responded to with adaptive strategy selection without having first to slowly learn about the differential performance of strategies. Thus, certain environmental characteristics might trigger certain strategies. Other characteristics might not directly trigger one particular heuristic or building block, but they can, as already suggested in the first chapter, at least exclude the use of certain building blocks. To adhere to the toolbox metaphor, a slotted-head screw prevents the use of a Phillips screwdriver. Analogously, the lack of feedback excludes constructing search orders that take the number of correct decisions per cue into account.

But not all environmental characteristics are easy to pick up or reveal themselves to the decision maker. This leads us to a third open question, closely related to the previous point: What does a decision maker need to know about the environment in order to be able to apply a certain building block? What are the necessary preconditions? Again, a simple example from the world of mechanic tools illustrates the point: For successfully using a wrench to unscrew the pedals from one's bicycle, one must know in which direction to turn it. Going back to the realm of decision strategies, let us use Take The Best as an example. Its decision rule only requires that one knows the direction in which a certain cue points. The stopping rule is even easier, requiring only to detect discriminations – which can potentially be more difficult with continuous cue values, especially in combination with pictorial cues, but not in the case of binary cues. Yet, when it comes to Take The Best's search rule, the task becomes less trivial. Take The Best presupposes that one knows which cues in a decision environment are valid and what their order by validity is. Even if one, in principle, wanted to start searching with the most valid piece of information and even if searching by validity was the most adaptive thing to do in a given environment, one will not be able to apply this search rule unless one knows this order. Or to turn the argument around: When one already knows what information to look up and which pieces of information are most valid, the decision is tremendously simplified. Consequently, Take The Best's success rests on work that is assumed to have been done beforehand, and how this work can be accomplished in simple ways was largely neglected. I therefore addressed this issue in the last chapter. I proposed adaptive ordering processes that lead to (relatively) accurate and very frugal decision making when combined with one-reason stopping and deciding. Indeed, experimental data supports the hypothesis that, in the online process of decision making, people use simple rules to order cues based on unconditional tallies of, for example, the number of correct decisions made by a cue. In terms of fit, these simple rules very clearly beat a rule that reorders cues by validity, whose fit was hardly better than that of a random model.

Results such as these might question the descriptive validity of some assumptions in the adaptive toolbox approach – possibly even those that are partly responsible for the most impressive results. For instance, the demonstration that one-reason decision making can beat complex information integration strategies, often used as normative standards in the field, in terms of accuracy is based on the assumption that search follows ecological validities. Yet such sobering findings are necessary contributions towards a more comprehensive understanding of also the limits to people's adaptivity. Of course, alternative sources of people's knowledge about the validity of cues are conceivable: They might have learned cue validities in circumstances other than decision-making tasks or deduced them from knowledge about causal relationships between cues and criterion. Information about important cues might have been transferred through social or evolutionary learning. These processes might also explain the origin of other building blocks and of knowledge about when to use them. But these potential mechanisms also need to be empirically tested.

In general, this dissertation tried to shed light on some of the still dim parts in the adaptive toolbox approach to decision making. But by shedding light on these, other dark areas were discovered. Thus, several issues will need further exploration, from simply testing some of the hypotheses derived in the first chapter to the more challenging questions about the origin of adaptive strategy selection even without learning from outcome feedback, which is demonstrated in the second chapter. Similarly, it will be an interesting challenge to investigate potential alternative sources of cue validity knowledge apart from learning from feedback in the online process of decision making. Only with continued empirical research efforts can the claim of psychological plausibility be turned into evidence, and the status of the adaptive toolbox as a theory of human judgment and decision making be advanced.