CHAPTER 2

Putting Adaptive Strategy Selection to a Test: Information Redundancy in Probabilistic Inferences

Introduction

Many of our everyday decisions are made under conditions of uncertainty. We cannot perfectly predict whether the route to work we choose will avoid the morning traffic jam, whether the journal will publish the manuscript we have submitted, or whether we will enjoy the movie we decide to rent. In all these cases, we have to rely on cues that are imperfectly correlated with the criterion we want to predict. Although they share this general characteristic, the decision problems we face differ on other dimensions. Sometimes it is difficult and expensive to gather information about one's decision alternatives, but at other times we have abundant information available for free. Sometimes many reasons speak for one alternative, whereas at other times different pieces of information contradict each other and point to different alternatives. The latter dimension, that is, the degree of information redundancy in a decision environment (measured via the correlations between cues), is what I will focus on in this chapter. This chapter will thus follow up on what has been said about information redundancy in the previous chapter (pp. 39-40), and explore in more depth which heuristics should be selected, and which are in fact selected in response to high versus low information redundancy in the decision environment.

Depending on the characteristics of the decision environment we encounter, we can respond in ways that increase the likelihood that our decision will be the right one. The approaches taken to explain such adaptive decision making have been very different in nature. Some authors have suggested general purpose, often computationally complex models that require that their parameters be adjusted to the particular situation. These include exemplar-based models (e.g., Juslin & Persson, 2002), linear integration models (e.g., Koehler, White & Grondin, 2003), and sequential sampling models (e.g., Lee & Cummins, 2004). Although the underlying assumptions of these models are quite different, they all share the property of free parameters that provide high flexibility in predicting the outcomes of different types of inference processes, but they also all have the drawback of creating potentially overly complex models.

In contrast, other authors have argued that people are equipped with a repertoire of decision strategies that often do not entail any free parameters (e.g., Gigerenzer et al., 1999; Payne et al., 1988, 1993; Svenson, 1979). In response to certain environmental structures, strategies are *selected* adaptively. The selection of strategies is often described as a process of trading strategies' costs (i.e., strategies' information processing demands) against strategies' benefits (i.e., strategies' accuracies; Beach & Mitchell, 1978; Payne et al., 1988, 1993).

How do these two general approaches differ in addressing the issue of adaptivity to the degree of redundancy? The influence of information redundancy has been addressed by the neo-Brunswikian "social judgment theory" (Brehmer & Joyce, 1988; Cooksey, 1996; Doherty & Kurz, 1996). Within this approach, multiple cue probability learning has been used extensively to study how people learn to estimate an object's criterion value based on several cues when making inferences repeatedly (Smedslund, 1955; for a short review, see Holzworth, 2001). The research closely follows Brunswik's (1955) ideas on the probabilistic nature of the information we rely on, and his claim for representative design that might involve interrelated cues. Multiple regression is able to capture uncertain cue–criterion correlations as well as correlations between cues, both integrated into the beta-weights of the regression equation. Consequently, research on multiple cue probability learning mainly models judgment by multiple regression.

Positive cue redundancy, defined as high positive correlations between cues, was found to be positively correlated with judgment accuracy (Naylor & Schenck, 1968) and speed of learning (Knowles, Hammond, Stewart & Summers, 1971). Results were, within the framework of multiple regression, mostly discussed in terms of people's ability to integrate not only cue–criterion but also cue–cue correlations into beta-weights. Armelius and Armelius (1974) doubt that people are capable of such an integration and report evidence that participants' beta-weights match cue-criterion correlations rather than the ecological beta-weights that take correlations between cues into account. A similar criticism comes from Schmitt and Dudycha (1975). This, however, did not stimulate a reconsideration of the use of multiple regression to model human judgment.

The methodological practice is problematic, though. Typically, the beta-weights of the regression model that best predicts participants' inferences are compared to the beta-weights of the best-performing regression model for solving the task (Brehmer, 1974). This method leaves open the possibility that people are using inference strategies that differ from linear regression models. Under conditions of high information redundancy, for example, ignoring information is often possible without much loss in predictive power, which might encourage the use of simple strategies that are more in reach of people's cognitive capacities. However, the adaptive use of such strategies might still lead to beta-weights in the fitted regression model that deviate from the beta-weights of the best-performing regression model for solving the task. This alternative explanation for the positive effects of redundancy has not been considered explicitly in the multiple cue probability learning literature, although some hints along these lines can be found in Brehmer (1974, p. 408). Hammond (1996) admits that the predominant use of the multiple regression technique for capturing human decision making might have prevented considering alternative, descriptively potentially more accurate models, including the possibility that people react to different degrees of information redundancy by using different inference strategies.

In the domain of preferential choices, Payne et al. (1988, 1993) have, in fact, argued that people are equipped with a repertoire of decision strategies and select strategies adaptively depending on the decision situation. Bettman et al. (1993) demonstrated that in a redundant environment with positively correlated attributes, the information search patterns participants displayed were more in line with the application of simple strategies that use only a subset of the available information. Negatively correlated attributes, in contrast, induced participants to process more information (but see also Johnson et al., 1989, for a failure to demonstrate adaptivity to correlational structure). But as said, these studies focus on preferential choices, which differ from inferential choices in that they depend on people's subjective evaluations, and no outside criterion exists to evaluate accuracy. Thus, it is an open question whether people react to information redundancy in the domain of inferences by changing their inference strategies, and if so, whether this change is adaptive.

The following studies will address this question. In a simulation study it will be analyzed how redundancy in probabilistic environments affects accuracy and frugality of different inference strategies. Based on the results I predict that simple one-reason inference heuristics will be best in predicting people's inferences under the condition of high redundancy of information. This prediction is tested in two experimental studies. Besides testing how adaptive people's decision processes are, individual differences in strategy application are examined more closely. Additionally, the second experiment explores whether learning about strategies' differential accuracy through outcome feedback is a necessary precondition for adaptive strategy selection. I conclude by discussing whether the experimental results support the *adaptive strategy selection hypothesis*, which states that people are equipped with several strategies and select strategies adaptively for different situations.

Simple heuristics for probabilistic inferences

Consider the inference problem of choosing between two potential oil fields the one with the higher quantity of oil. For this inference, different tests could be carried out – chemical analyses, for instance, to determine the content of the organic matter in the bedrock (Mobil Oil, 1997). But which and how many tests should be considered, and how should the test results be used to make an inference? Gigerenzer and Goldstein (1996, 1999) have proposed that people's inferences are often based on simple heuristics. Among the proposed heuristics, Take The Best (TTB), which has been introduced in the first chapter, has received the most attention so far, both in simulation and experimental studies. This one-reason decision-making heuristic is a non-compensatory inference strategy, because a cue cannot be outweighed by any combination of less valid cues, in contrast to a compensatory strategy, which integrates cue values. In simulations, TTB proved to be highly accurate compared to compensatory strategies across a wide range of real-world environments (Czerlinski et al., 1999). Its simplicity and accuracy make TTB a psychologically plausible model of human inference.

What are the conditions under which people actually apply such a non-compensatory strategy? Recent studies, reviewed in the first chapter, have addressed this question (Bröder, 2000, 2003; Bröder & Schiffer, 2003; Newell & Shanks, 2003; Newell et al., 2003, Rieskamp & Hoffrage, 1999, 2003). All this work relies on the adaptive strategy selection hypothesis, which states that simple heuristics should predict inferences best when they outperform alternative, more complex strategies. A strategy's performance depends on its accuracy and on its application costs. A strategy's accuracy can be defined as the percentage of correct predictions the strategy allows, and a strategy's application costs as the amount of information the strategy requires as well as the effort necessary for processing the information. For parsimony, processing costs are ignored, and the focus is solely on the required amount of information of a strategy, that is, on a strategy's frugality, as the main source of a strategy's costs. Also for parsimony, a strategy's overall performance is defined as the monetary payoff the strategy achieves.

The adaptive strategy selection hypothesis has received support in the studies cited above: As seen in Chapter 1, simple heuristics predict inferences well when the costs for applying compensatory strategies are high. For instance, simple heuristics reached a high fit in predicting inferential choices when participants had to pay for the information they acquired (Bröder, 2000; Newell & Shanks, 2003, Newell et al., 2003), when they made inferences under high time pressure (Rieskamp & Hoffrage, 1999), or when information had to be retrieved from memory (Bröder & Schiffer, 2003).

To influence strategies' overall performance, most of the reported studies have thus varied conditions that affect strategies' costs rather than strategies' accuracy. In particular, the influence of information redundancy on strategies' accuracy has gained very little attention, but it is very intuitive that information redundancy should have a strong effect. Consider, for instance, an inference situation with three cues. When information redundancy is high, cues are highly correlated with each other so that relying only on the information of one single cue should cause no harm. In contrast, in a situation of low information redundancy in which, for example, the most valid cue is not correlated with the other cues, these cues provide additional information. Consequently, it appears intuitively plausible that the best thing to do is to integrate this additional information. The Simulation Study examines whether this intuitive expectation is correct.

Simulation Study

The goal of the Simulation Study is to examine how information redundancy in decision environments affects the accuracy and frugality of various inference strategies, by means of computer simulations. The inference task can, in extension of what has been presented in the first chapter, be conceptualized as follows: An environment consists of a population of N objects, of which each object j is characterized by a criterion value x_j. The task is to predict for all possible paired comparisons which object has the larger criterion value. Each object is described by a set of M cues. Each cue i can take a positive or negative cue value c_i (i.e. 1 and 0). The cue difference g_i denotes the difference of the cue value of the first alternative A minus the cue value of the second alternative B. Information redundancy occurs when the cues are correlated with each other. First, I will describe three strategies that can be used to solve this task and that differ in computational complexity and information demands. Second, I will examine how the performance of these strategies is affected by information redundancy.

Strategies for probabilistic inferences

The first strategy that will be considered is *TTB*, described earlier. It is a non-compensatory strategy that does not integrate any information. The second strategy in the competition is *Naïve Bayes* (NB). This strategy is used to represent the class of compensatory strategies, since it integrates the information of all available cues. Some authors have argued that NB should be regarded as the "rational" model for probabilistic inferences (Lee & Cummins, 2004). In fact, Martignon and Laskey (1999) could demonstrate that NB outperforms TTB in making probabilistic inferences across a large number of real environments. Bröder (2000), Newell and Shanks (2003), and Newell et al. (2003) have used NB in their experiments to determine which alternative was considered as the correct choice, and provided feedback and payoff for participants accordingly. NB estimates the probability p that the first alternative has a larger criterion value than the second alternative. Its prediction can be determined by the posterior odds that A has a larger criterion value than B given a particular cue profile. Transformed onto a log-odds scale, the posterior odds can be computed through adding the log-odds for each cue (derived from the cue's validity) multiplied by the encountered cue difference:

$$\ln\left(\frac{p_k(x_A > x_B)}{1 - p_k(x_A > x_B)}\right) = \ln\left(\frac{v_1}{1-v_1}\right) g_1 + \ldots + \ln\left(\frac{v_i}{1-v_i}\right) g_i + \ldots + \ln\left(\frac{v_M}{1-v_M}\right) g_M ,$$

where k is a particular paired comparison, and v_i is the validity of cue i. When the predicted score is larger than 0, then the estimated probability that the first alternative has a larger criterion value compared to the second alternative is greater than .50, so that alternative A should be selected (and vice versa). NB integrates the information of all available cues, but it ignores correlations between cues, making the simplifying assumption that cues are independent.

Considering how NB determines its decision, it might be too complicated a model to be psychologically plausible. However, a compensatory strategy does not need to be computationally demanding. The third strategy, which I call *Take Two,* represents a compensatory strategy that is simple to apply. It thus builds a bridge between TTB and NB: Like TTB, it searches for cues in the order of their validity. But in contrast to TTB, it only stops search when two cues that favor the same alternative have been found; that alternative is then selected regardless of whether during search another cue was found that favored the other alternative. If Take Two does not find two cues that favor the same alternative, it selects the alternative that is favored by the cue with the highest validity. The strategy is based on the idea that people sometimes are not comfortable basing their decision on one single cue but nevertheless might want to limit their information search. Take Two satisfies these motivations. It has the interesting property that it can lead to intransitive choices.[13]

Testing the strategies

The main goal of the Simulation Study is to identify which strategies perform well under high information redundancy, and which perform well under low redundancy. To test the strategies' accuracy I focus on generalizability as the main evaluation criterion (Pitt & Myung, 2002). A strategy's generalizability can be defined as its ability to solve inferences in a new environment, without adapting the strategies' parameters again. Therefore, a cross-validation study was performed where half of the objects of an environment were used as a calibration sample for determining the validities of the cues and their rank order, and the other half were used as a validation sample for testing the strategies' generalizability. An interesting side question therefore is whether NB – which relies on the exact ecological cue validities – proves to be less robust when applied to a new environment than TTB or Take Two, which merely have to derive the cues' rank order based on their ecological validities.

In addition to accuracy, the strategies' frugality will be evaluated, defined as the average number of cues acquired for making an inference. Whereas TTB and Take Two define how they search for information and when information search stops, this is not clear for NB. At first glance, one expects that all available cues will be required. However, even

---

[13] To illustrate Take Two's intransitivity, imagine a set of three alternatives, A, B, and C. Each alternative is described by six cue values, which are in the order of their validity $c_1 = 1$, $c_2 = 1$, $c_3 = 0$, $c_4 = 0$, $c_5 = 0$, $c_6 = 1$ for alternative A, $c_1 = 1$, $c_2 = 0$, $c_3 = 1$, $c_4 = 0$, $c_5 = 0$, $c_6 = 0$ for alternative B, and $c_1 = 0$, $c_2 = 1$, $c_3 = 0$, $c_4 = 1$, $c_5 = 1$, $c_6 = 0$ for alternative C. When comparing alternative A and B Take Two selects A, because it is favored by the second and sixth cue. When comparing alternative B and C, Take Two selects B, because it is favored by the first and third cue. However, when comparing alternative A and C, Take Two selects C, because it is favored by the fourth and fifth cue, leading to an intransitive circle.

for NB limited information search is, in principle, possible. Suppose the five most valid of six cues favor one alternative: Searching for the sixth cue is not necessary because it would not change the prediction of NB. Thus, search can be limited by assuming that NB also searches for cues in the order of their validities and stops search when additional cues cannot change a decision based on the cues acquired so far. It is of course questionable whether this search process is psychologically plausible, since it requires hypothetically determining a decision after each acquired cue, and checking whether it is, hypothetically, still possible to arrive at a different final decision. Nevertheless, this limited search process is assumed for NB, as it reduces the necessary amount of information. This leads to a more demanding competition between the strategies regarding their frugality and enables a more conservative test of TTB's apparent advantage regarding frugality compared to NB.

Method

Strategies' accuracy and frugality were tested in environments with either high or low information redundancy. In addition, the distribution of the cue validities was varied: In the high dispersion condition, validities varied widely, whereas in the low dispersion condition validities varied only moderately. The dispersion of cue validities should also influence strategies' accuracy: Intuitively, relying on the most valid cues when they have a much higher validity compared to remaining cues appears more reasonable than when their validities do not differ substantially. The two factors, information redundancy and validity dispersion, were crossed, providing four environment conditions.

Table 2.1: Average cue validities and average correlation between cues in the four kinds of decision environments used in the Simulation Study.

	High information redundancy		Low information redundancy	
	High dispersion	Low dispersion	High dispersion	Low dispersion
Validities of cues				
First cue	.89	.82	.89	.81
Second cue	.82	.78	.82	.77
Third cue	.76	.74	.75	.73
Fourth cue	.69	.70	.68	.69
Fifth cue	.62	.66	.61	.65
Sixth cue	.56	.62	.54	.61
Average correlation between cues	$r = .51$	$r = .51$	$r = .01$	$r = .01$

In more detail, 500 artificial environments, each consisting of 50 objects and 6 cues, were created for each of the four conditions. Every cue had 25 positive cue values and 25 negative cue values. Environments were constructed by a computer program that first randomly distributed cue values to the objects, and thereafter interchanged cue values until redundancy and cue validity requirements were met. In the high-redundancy environments, cues correlated substantially with each other; in the low-redundancy environments, cues on average did not correlate with each other (see Table 2.1).

Results

Accuracy of strategies

Table 2.2 lists the average proportion of correct inferences made by the different strategies in the validation sample and the calibration sample, under high and low cue redundancy, and low and high validity dispersion. Strategies' accuracy relative to each other is influenced by the environment: In the calibration set, TTB and NB achieved very similar accuracy in the high-redundancy environments, whereas NB did on average better than TTB in the low-redundancy environments. The lowest accuracy in the high redundancy conditions was achieved by Take Two, which in the low redundancy conditions achieves intermediate accuracy. Overall, accuracy differences between the strategies were small.

Table 2.2: Proportion of correct inferences achieved by the three strategies in four kinds of decision environments in data fitting (calibration) and cross-validation.

		High information redundancy		Low information redundancy	
		High dispersion of cue validities	Low dispersion of cue validities	High dispersion of cue validities	Low dispersion of cue validities
NB	Calibration	.71	.70	.87	.86
	Validation	.70	.68	.84	.82
Take Two	Calibration	.69	.68	.85	.84
	Validation	.68	.68	.83	.82
TTB	Calibration	.72	.71	.85	.82
	Validation	.70	.68	.82	.74

The more important question, however, is how the strategies perform when making predictions for independent data, that is, their generalizability. The loss in accuracy from the calibration sample to the validation sample varied between strategies and environments. NB, which multiplies cue values by weights derived from their validities, was not less robust than TTB or Take Two, with a maximum loss of 4 percentage points in the low redundancy and low validity dispersion condition. Take Two was the most robust strategy in the competition, with a maximum loss of 2 percentage points in both low redundancy conditions. TTB experienced a loss in accuracy of 8 percentage points in the low redundancy and low validity dispersion condition, which is the largest loss that was observed in this simulation. In the other three conditions, its robustness was similar to NB. When comparing the absolute accuracy levels achieved in the validation environments, TTB, the most frugal strategy in the competition, demonstrated high accuracy especially in the high redundancy conditions, where it performed as good or even better compared to NB and Take Two. Yet even under low redundancy conditions, when dispersion of cue validities was high, TTB performed on a very similar level to NB. The only situation in which TTB suffered a clear loss in accuracy, both when comparing calibration and validation sample and in comparison to other strategies, was when cues were low in redundancy *and* had similar validities. Then, its accuracy dropped by eight percentage points to 74%, lagging eight percentage points behind Take Two, which in turn is as good as the computationally more complex NB.

Frugality of strategies

Table 2.3 shows the average number of cues acquired by the strategies to arrive at a decision. Since the fitting process does not affect strategies' frugality, frugality was averaged across the calibration and validation sample. As described above, limited information search was also assumed for NB. Even with this additional assumption, making the competition between the strategies more demanding, TTB clearly required less information than the two compensatory strategies. NB, despite searching until a decision based on the cues encountered so far can no longer be overruled, was slightly more frugal than Take Two. However, NB's advantage over Take Two in terms of frugality can be attributed to the restricted number of six cues. If the number of available cues were increased, Take Two would become more frugal relative to NB, as, regardless of how many cues there are to come, it will stop once two cues that support one alternative are encountered.

Table 2.3: Average number of cues looked up by the three strategies in the four kinds of decision environments.

	High information redundancy		Low information redundancy	
	High dispersion of cue validities	Low dispersion of cue validities	High dispersion of cue validities	Low dispersion of cue validities
NB	3.4	3.9	3.5	4.0
Take Two	3.9	4.0	4.2	4.3
TTB	2.3	2.3	2.0	1.9

Discussion

The main goal of the Simulation Study was to identify strategies that perform well under high and low information redundancy. The prediction that one-reason decision making, as applied by TTB, performs very accurately in high-redundancy environments whereas compensatory strategies are more accurate in low-redundancy environments generally holds. Surprisingly, however, when cue validities vary substantially TTB still reaches a relatively high accuracy level. The only situation in which TTB is outperformed by the compensatory strategies is when cues with similar validities offer non-redundant information.

An interesting side question was whether a more complex strategy like NB would behave less robustly (Myung & Pitt, 2002) and suffer a higher loss in accuracy from calibration to validation. This was not the case, and losses in accuracy were mostly within a similar (and low) range for all strategies, with the exception of TTB that suffered a large loss in accuracy when both information redundancy and dispersion of cue validities was low.

With regard to the generality of these results, one might argue that TTB would suffer even more in environments in which cues are not only non-redundant, that is, not correlated, as in the low-redundancy condition of the Simulation Study, but are negatively correlated with each other, so that cues favor different alternatives, thereby creating conflict (Shanteau & Thomas, 2000). Negative correlations imply that trade-offs between different cues have to be made. Several authors have argued that such trade-offs are a fundamental aspect of choice in any domain (Fasolo et al., 2004; Stillwell, Seaver & Edwards, 1981), and that "if there were no conflict, there would be no choice" (Hogarth, 1980, p.59), but these authors refer to preferential choices. For inferential choices, in contrast to preferential choices, an objective outside criterion exists on which to evaluate the accuracy of an inference. Therefore, only cues that are positively correlated with the criterion will be considered (i.e., cues with a validity greater than .50). The constraint that all cues are positively correlated with the criterion also increases the average correlation between cues. Inference situations in which cues are negatively correlated with each other are therefore rare. For preferential choice, the situation is different: The decision maker can assign positive and negative *subjective* weights (representing importance) to the attributes of alternatives. Preferential choices can therefore often be characterized by conflict in terms of negative inter-attribute correlations (e.g., between prize and quality of a product).

What can the results of the Simulation Study contribute to predictions about the strategies people should select in response to information redundancy? It has been shown that the simplest strategy under consideration – TTB – achieves high accuracy under conditions of high information redundancy. Therefore, no trade-off between accuracy and costs has to be made: Taking into account the heuristic's high frugality, TTB will achieve a higher net payoff when costs for information are incurred, and therefore it will be the most adaptive strategy to use when information search becomes costly. Only in the condition of low information redundancy and low dispersion of cue validities is TTB outperformed by compensatory strategies. In this situation, NB and Take Two achieve a higher accuracy.

Experiment 1

According to the adaptive strategy selection hypothesis people select strategies for solving an inference problem according to their performances, which depend on the environment in which they are applied. This hypothesis is tested in Experiment 1, which focuses in particular on the question of whether people respond adaptively to information redundancy. The Simulation Study showed that information redundancy affects strategies' accuracies, such that the simple heuristic TTB performs well in environments with high information redundancy. In contrast, under low information redundancy, in particular with low dispersion of cue validities, TTB was outperformed by compensatory strategies. From these results it can be predicted that people select TTB for making their inferences in situations of high information redundancy and select compensatory strategies in situations of low information redundancy. To represent the class of compensatory strategies, Take Two (as a computationally simple heuristic) and NB were selected (as a computationally more complex but usually also slightly more accurate strategy on average).

However, the characteristics of an environment, in particular information redundancy, are often not obvious, so that people need to explore the environment before an appropriate strategy can be selected. Participants in Experiment 1 were therefore provided with a learning phase in which the environment could be explored without having to pay for information. After the learning phase, participants were confronted with the decision phase, which involved information search costs. The Simulation Study showed that another main advantage of TTB is its frugality compared to the compensatory strategies. Thus, when search costs are introduced, TTB's performance in terms of payoff increases relative to compensatory strategies. Consequently, people should limit their information search when search costs are incurred and should more often select TTB to make their inferences.

Furthermore, as a side question, I was interested in potential individual differences in responding to information redundancy. A central tenet of cost–benefit theories of decision-strategy selection (e.g., Beach & Mitchell, 1978; Payne et al., 1993) is that the payoff of a strategy determines its use in a given environment. This assumption leaves little room for individual differences. But it is known from research on preferential choice that negative inter-attribute correlations increase the likelihood of non-dominated options, thus making choices difficult (Fasolo, McClelland & Todd, 2003; Tversky & Shafir, 1992), and possibly leading to less satisfaction and confidence of the chooser (Jacoby, Speller & Berning, 1974; Malhotra, 1982). People might thus seek to avoid explicit trade-offs when facing an environment with negatively correlated attributes, as suggested by Hogarth (1980), and one way to do so is to use non-compensatory decision strategies (Dhar, 1996). If this argument holds also for the inference domain, the prediction would be the opposite of the first hypothesis, namely, more TTB use will be observed in the low-redundancy environment.

Evidence gathered by Bettman et al. (1993), again in the preference domain, favors the hypothesis derived from cost–benefit analyses, but one can expect differences between people. One dimension of individual differences potentially implicated here is the "need for cognitive closure," defined as a desire for "*an* answer on a given topic, *any* answer, as compared to confusion and ambiguity" (Kruglanski, 1990, p. 337, emphasis in original). Need for cognitive closure is assumed both to vary as a function of the situation and to represent a dimension of stable individual differences (Webster & Kruglanski, 1994). Especially one subscale of the Need for Cognitive Closure questionnaire (Kruglanski, Webster & Klem, 1993) seems relevant here: "Discomfort with Ambiguity." It refers to the affective discomfort experienced by ambiguity, that is, the absence of closure (see, e.g., item 31: "It is annoying to listen to someone who cannot seem to make up his or her mind."). People who experience ambiguity as aversive might therefore use one-reason decision making as a means to avoid ambiguity-causing conflict of information.

In short, my hypotheses are as follows:

Hypothesis 1: The proportion of inferences predicted by TTB will be larger compared to compensatory strategies when information search is costly than when no direct costs for information are incurred.

Hypothesis 2: The proportion of inferences predicted by TTB will be larger compared to compensatory strategies in a high information-redundancy condition than in a low information-redundancy condition.

Hypothesis 3: In the low-redundancy environment, high discomfort with ambiguity, expressed on the Need for Cognitive Closure Scale (Kruglanski et al., 1993), is associated with a higher proportion of inferences predicted by TTB.

These three hypotheses were tested in Experiment 1.

Method

Participants

Forty participants (22 female, 18 male; average age 24 years), mostly students from various departments of the Free University of Berlin, took part in the study. They received performance-contingent payment for their participation of on average 10.32 €, ranging from 1.80 to 18.82 €. The computerized task, which was conducted in individual sessions, lasted approximately 1 hour. After completing this task, they filled out the German version of the Need for Cognitive Closure questionnaire (Deutsch, Bohner, Kemmelmeier & Erb, 2003), for which they received an additional 2.00 €.

Procedure

Participants were asked to imagine they were geologists hired by an oil-mining company. The instructions explained that their task was to decide at which of two potential drilling sites, labeled X and Y, more oil would be found based on various test results. Participants were told how many inferences they had to make, how much they would earn for a correct inference, and how much they had to pay for a wrong inference. They could conduct six different tests (“Chemical analysis,” “Geophones,” “Gravimetry,” “Ground water analysis,” “Microscopic analysis,” “Seismic analysis”). Tests could be selected successively by clicking on corresponding icons on the computer screen (see Figure 2.1). Each test’s validity was indicated below its icon and it was made clear which of the two possible test outcomes would favor high oil occurrence. When a test was selected, the test results were revealed for both drilling sites and remained visible until a decision was made. The positions at which the tests were displayed at the top of the screen were the same throughout all decisions for one participant but varied randomly across participants. The order in which test results appeared on the lower part of the screen was determined by the order in which the tests had been conducted. After a decision had been made, outcome feedback was provided by either a green “correct” box or a red “wrong” box.

Figure 2.1: Screenshot of the computerized task participants faced in Experiments 1 and 2 (taken from the decision phase).

The experimental design had two factors: phase (within subjects; learning phase vs. decision phase) and environment (between subjects; high redundancy vs. low redundancy). Overall, participants made 192 inferences without any time constraint. The initial learning phase and the final decision phase each consisted of 96 decisions, containing three repetitions of blocks of 32 items. Within each block, items were randomly ordered. On which side on the screen –

left or right – the correct alternative was presented was also determined randomly. Whereas in the learning phase participants could look up information for free, in the decision phase they had to pay acquisition costs for each test they conducted.

The environment was constructed such that in the high-redundancy-environment condition the average inter-cue correlation between the six available cues was $r = .50$ (with a minimum correlation between two cues of $r = .35$) compared to the low-redundancy-environment condition with an average inter-cue correlation of $r = -.15$ (with a maximum correlation between two cues of $r = -.02$). Thus, under low redundancy the cues not only provided additional valid information, but also revealed pieces of information that often were in conflict with each other. The cues' validities of 0.83, 0.78, 0.72, 0.67, 0.61, and 0.56 were the same for both environments (and all cues had the same discrimination rate of 0.56). In the high-redundancy environment, TTB and NB had the same accuracy of 78%, whereas in the low-redundancy environment, TTB made 66% correct decisions compared to NB with 91% correct decisions. As NB and Take Two always predicted the same inference outcomes, NB stands for the performance of both compensatory strategies. These accuracy levels match the results of the Simulation Study: TTB performed as well as compensatory strategies under high information redundancy but was outperformed under low information redundancy. Due to the redundancy between cues, the overlap in strategies' predictions is rather large in the high-redundancy environment, with identical predictions between TTB and the compensatory strategies of 94%. The corresponding number for the low-redundancy environment is 63%. Thus, in both environment conditions the selection of a compensatory strategy led to accuracy superior or equivalent to that of TTB – it is higher frugality that makes TTB perform better than compensatory strategies in terms of payoff in the high-redundancy environment once costs for cues are introduced.

Participants' payoff, presented on the computer screen, was expressed in an experimental currency called "petros." Ten thousand petros corresponded to 1.00 €. Participants earned 0.20 € for a correct decision and paid 0.20 € for a wrong decision. In the decision phase, they paid 0.03 € for acquiring a test result. Thus, in the learning phase, strategies' payoffs were solely determined by their accuracies, leading to equal payoffs for TTB and NB of 10.80 € in the high-redundancy condition, and 6.00 € for TTB and 15.60 € for NB in the low-redundancy condition. The information costs introduced in the decision phase, however, changed strategies' payoffs. In the decision phase of the high-redundancy environment the consistent application of TTB led to a payoff of 4.77 € compared to NB with a payoff of 1.62 €. In the low-redundancy environment performance was reversed, so that the consistent application of TTB led to a payoff of 1.41 € compared to NB with a payoff of 4.44 €. According to Hypothesis 1, participants should limit their information search, and more frequently apply TTB, in the decision phase compared to the learning phase due to search costs. According to Hypothesis 2, TTB should be better in predicting the inference process in the high-redundancy condition, in which TTB outperforms NB in terms of payoff compared to the low-redundancy condition, in which a more extensive information search is worthwhile as NB outperforms TTB.

Results

First, participants' inferences were analyzed by examining participants' overall performance. Second, both adaptivity hypotheses were tested by examining how well TTB predicted the inference outcomes and the information search patterns compared to the compensatory strategies.

Participants' performance

An analysis of variance with payoff as the dependent variable, phase (learning vs. decision phase) as a within-subject factor, and environment (high vs. low redundancy) as a between-subjects was conducted. The average payoff of 9.97 € was higher in the learning phase compared to the average payoff of 2.58 € in the decision phase (see Table 2.4), $F(1, 38) = 268.35$, $p = .001$, $\eta^2 = 0.88$, which is not surprising since the decision phase, contrary to the learning phase, involved search costs. In the low-redundancy environment, participants received on average higher payoffs (14.13 €) than participants in the high-redundancy environment (10.96 €), $F(1, 38) = 9.10$, $p = .005$, $\eta^2 = 0.19$. However, this effect only holds for the learning phase due to an interaction effect: Participants in the low-redundancy environment received higher payoffs in the learning phase, with on average 11.42 € compared to 8.52 € in the high-redundancy condition, but not in the decision phase (2.71 vs. 2.44 €), $F(1, 38) = 8.48$, $p = .006$, $\eta^2 = 0.18$. This is again not surprising as in the low-redundancy environment the predicted application of a compensatory strategy in the learning phase leads to a higher payoff than application of any strategy in the high-redundancy environment.

Table 2.4: Accuracy, frugality, and payoff achieved by participants in the high- and low-redundancy conditions of Experiment 1, reported separately for learning and decision phases.

	High information redundancy		Low information redundancy	
	Learning phase	Decision phase	Learning phase	Decision phase
Mean percentage correct (*SD*)	72% (*4.3*)	71% (*6.2*)	80% (*8.6*)	76% (*6.3*)
Mean number of cues looked up (*SD*)	4.5 (*1.5*)	2.0 (*0.7*)	3.7 (*1.3*)	2.6 (*0.7*)
Mean payoff (*SD*)	8.52 € (*1.7*)	2.44 € (*1.8*)	11.42 € (*3.3*)	2.71 € (*1.5*)

The introduction of search costs in the decision phase changed the situation. Now, participants' payoffs were a function of how accurately they made their inferences and how frugal they were in their information search. Although participants on average still achieved higher accuracy in the low-redundancy environment, with on average 76% correct decisions, compared to 71% in the high-redundancy environment, $t(38) = 2.51$, $p = .008$ (one-tailed), $d = 0.80$, they were less frugal. Participants in the low-redundancy environment looked up 2.6 cues per trial on average, while participants in the highly redundant environment looked up 2.0 cues on average, $t(38) = 2.48$, $p = .009$ (one-tailed), $d = 0.78$. This explains why the two groups do not differ in their payoffs in the decision phase. In sum, these results support the adaptivity assumption. Although environmental structure differed, participants performed equally well in terms of payoff in the decision phase of both conditions.

Testing the strategies

It is first tested how well TTB and the compensatory strategies (NB and Take Two, which make the same outcome predictions) predicted participants' inferences, and second how well they predicted participants' information search (where NB and Take Two are tested separately, as they predict different stopping points).

In the learning phase TTB predicted on average 89% of the inferences in the high-redundancy condition compared to 69% in the low-redundancy condition. In comparison, NB predicted 92% of the inferences in the high-redundancy condition and 84% of the inferences in the low-redundancy condition. Thus, the compensatory strategy predicted more inferences in both environments.

In the decision phase, TTB predicted 87% of the inferences in the high-redundancy condition compared to 77% in the low-redundancy condition. NB predicted 85% of the inferences in the high-redundancy condition compared to 79% in the low-redundancy condition. Now, as predicted, TTB is the strategy that predicts the most inferences in the high-redundancy condition whereas compensatory strategies predict more inferences in the low-redundancy condition. However, due to the large overlap in strategies' predictions the differences in fit are small and do not allow a conclusive comparison of the strategies on the basis of the inference outcomes. In this respect, information search becomes the crucial indicator of which strategies were applied.

Two aspects of information search were analyzed: search order and stopping of search. Whereas in the learning phase participants were able to explore the environment without having to pay for information, so that frugality should not be a main concern, frugality became crucial in the decision phase: Adaptive inference processes lead one to limit information search and to search for information in the order of cue validities.

In which order did participants search for information? If participants select TTB to make their inferences they should have searched for cues according to their validity. However, as argued above, even when using a compensatory strategy it is best to search for cues in the order of their validity to limit information search. In the high-redundancy condition, participants searched in the order of cue validities in 47% of all inferences in the learning phase compared to 73% of the inferences in the decision phase, $t(19) = 3.21$, $p = .002$ (one-tailed), $d = 0.72$. Similarly, in the low-redundancy condition, participants' search followed validity in 57% of the inferences in the learning phase compared to 72% of the inferences in the decision phase, $t(19) = 2.82$, $p = .005$ (one-tailed), $d = 0.62$.

These results demonstrate that when search costs are introduced participants respond adaptively by more frequently searching for cues in the order of their validity. However, the search order does not discriminate between different strategies. For this purpose, the crucial test is when information search stops. If participants selected TTB to make an inference, information search should stop when the first discriminating cue is found. I first determined the proportion of inferences in which search stopped before any discriminating cue was found. Since in such a situation no alternative is favored by the available information, this type of search indicates that participants made a guess about which alternative to choose. In the learning phase, search stopped in only 2% of inferences before a discriminating cue was found compared to 11% in the decision phase.

To test whether participants stopped search in accordance with TTB or rather in line with a compensatory strategy, I determined for all inferences in which participants did not guess whether search stopped when one discriminating cue was found (as predicted by TTB) or continued (as predicted by compensatory strategies). First, consistent with Hypothesis 1, the introduction of search costs led to limited information search. For both conditions, adherence to TTB's stopping rule increased substantially from the learning to the decision phase: In the learning phase of the high-redundancy condition, participants stopped their search in 23% of the non-guessing inferences when one discriminating cue was found, compared to 77% in the decision phase, $t(19) = 6.03$, $p = .001$ (one-tailed), $d = 1.35$ (see also Figure 2.2). In the learning phase of the low-redundancy condition, participants stopped their search in 26% of the non-guessing inferences when one discriminating cue was found, compared to 44% in the decision phase, $t(19) = 3.72$, $p = .001$ (one-tailed), $d = 0.83$.

To test Hypothesis 2, I compared stopping behavior in the decision phases of the two redundancy conditions. As predicted, the proportion of inferences in which search stopped when one discriminating cue was found, consistent with TTB, was higher in the high-redundancy condition than in the low-redundancy condition (77 versus 44%), $t(38) = 3.37$, $p = .001$ (one-tailed), $d = 1.07$. Interestingly, this effect can already be found in the first block of the decision phase, block 4, with TTB-consistent stopping in 76% of trials in the high-redundancy condition, and 43% in the low-redundancy condition. Figure 2.2 shows for both conditions the proportion of non-guessing trials in which search stopped in accordance with TTB's stopping rule, that is, as soon as one discriminating cue was found, and the complement proportion of trials in which search continued beyond a first discriminating cue (across six blocks of trials).

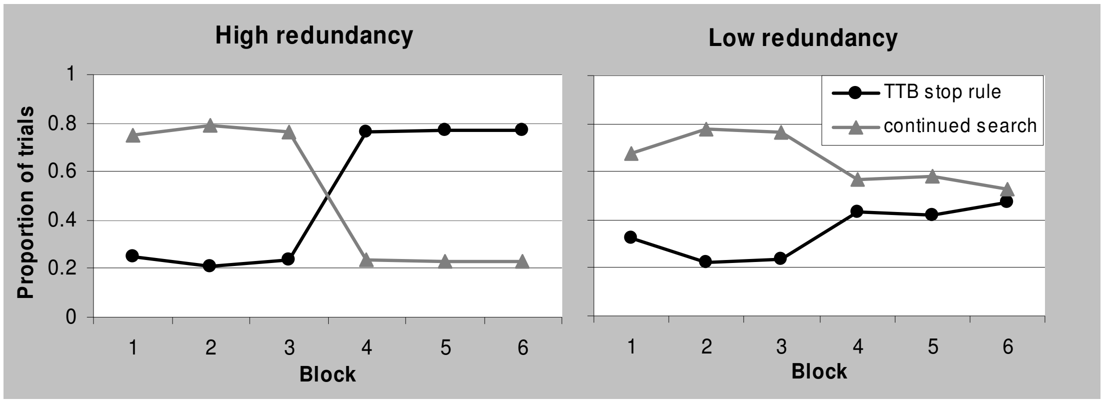

Figure 2.2: Proportion of non-guessing trials in which search stopped in accordance with TTB (i.e., when one discriminating cue was found) compared to the complement proportion of instances in which search continued beyond a first discriminating cue, across the six blocks of trials in Experiment 1.

Thus, whereas in the high-redundancy condition search predominantly stopped after one discriminating cue was found, as predicted by TTB, in the low-redundancy condition search mostly continued, as predicted by compensatory strategies. But which of the two compensatory strategies predicts the information search pattern better? Again for all non-guessing cases I analyzed whether participants stopped search as predicted by Take Two or NB.[14] In the decision phase of the high-redundancy condition search stopped in accordance with Take Two in 20% of all non-guessing inferences compared to 16% in accordance with NB, $t(19) = 2.49$, $p = .011$ (one-tailed), $d = 0.57$. In the decision phase of the low-redundancy condition, participants stopped search as predicted by Take Two in 31% of the non-guessing inferences compared to 22% as predicted by NB, $t(19) = 3.17$, $p = .003$ (one-tailed), $d = 0.86$. Thus, Take Two in both conditions predicts better than NB when participants stop their information search. Also, Take Two predicts information search slightly better in the low-redundancy compared to the high-redundancy condition (31% compared to 20%), $t(37) = 1.44$, $p =.079$ (one-tailed), $d = 0.46$. In contrast, exhaustive search was extremely rare in the decision phase: Participants looked up all cues in only 2% of the non-guessing inferences in the high-redundancy condition and in 5% in the low-redundancy condition.

Is the information search pattern a valid indictor of the strategy that is used?

As said, the overlap between strategies makes it impossible to conclude from the decision outcomes which decision strategy has been used predominantly. I therefore resorted to

[14] The overlap between the predictions of the stopping behavior by Take Two and NB is very high: They predict the same stopping point in 82% of the cases. This explains why even small differences result in a substantial effect: Although small, the differences almost exclusively favor Take Two over NB, suggesting that Take Two clearly makes more correct predictions beyond those that the two strategies have in common.

process measures, and among these, mainly to stopping behavior. But how valid an indicator is stopping for the decision rule used?

Let us first consider stopping when one, and only one, discriminating cue has been found. When such "one-reason stopping" is applied, the only reasonable decision strategy is to rely on that cue. Indeed, participants decided 99% of these cases in the decision phase in accordance with this first (and only discriminating) cue, regardless of environment condition. Looking up more information, however, opens new possibilities with regard to which decision rule to apply. In principle, one could still rely on the first discriminating cue. However, this did not happen. I selected all cases in which search continued until the point predicted by Take Two in the decision phase (or beyond). In the high-redundancy environment, decisions were made in accordance with Take Two in 97% of these cases. While this could still be partially explained by the high overlap between the predictions of the strategies, in the low-redundancy environment, strategies more often made opposing predictions. Nevertheless, the percentage of decisions in accordance with Take Two was not diminished: Of the cases in which the necessary information had been looked up, 98% of decisions were in line with Take Two. Even when focusing only on the subset of critical profiles, for which TTB and Take Two *differ* in their predictions (i.e. after one discriminating cue, two more are found that point into the other direction), results hardly change. Participants decided in favor of the alternative to which the two cues pointed in 96% of these trials, leaving only 4% being in favor of the first discriminating cue. Taken together, these results indicate that if people search beyond the first discriminating cue, they also use the gathered information in a way that can overrule the first discriminating cue. Thus, the conclusion that stopping is a valid indicator of the kind of decision strategy that is being used – non-compensatory or compensatory – is supported. More generally, the results demonstrate that stopping and decision rules are not independent of each other.

Information search and discomfort with ambiguity

Hypothesis 3 predicted that discomfort with ambiguity as expressed on the Need for Cognitive Closure Scale would be associated with higher levels of TTB use in the decision phase if the environment consisted of cues that often contradicted each other (as in the low-redundancy condition). I therefore correlated the proportion of TTB-consistent stopping in the decision phase of the experiment (blocks 4–6) with the participants' average score on the discomfort with ambiguity subscale. As expected, no correlation was found in the high-redundancy environment, r = .07, p =.383 (one-tailed). But, contrary to Hypothesis 3, the correlation was also small, and even negative, for participants in the low-redundancy environment, r = –.28, p = .880 (one-tailed).

Discussion

The results of Experiment 1 provide evidence for the first two hypotheses. Participants seemed to respond to information costs and degree of information redundancy by using decision strategies that are adaptive in the encountered environment conditions. When information became costly, participants in both groups mostly searched cues in an order corresponding to their validities. Participants in the high redundancy environment reacted to the information costs mainly by one-reason stopping and deciding. Participants in the low-redundancy environment also became more frugal when information costs were introduced. But unlike participants in the redundant environment, they mostly continued search beyond a first discriminating cue. Among the cases of continued search, Take Two was better than NB in predicting when search was stopped. Nevertheless, even in the low-redundancy environment, TTB's stopping rule achieved a higher fit than any of the other strategies that were considered. This shows that it is difficult for compensatory strategies that are specified not only in terms of decision outcomes but also processes to compete with TTB in predicting participants' decision behavior.

The results have interesting implications for the interpretation and validity of process measures, in particular stopping behavior. When search is stopped after the first discriminating cue, participants almost always choose the option the cue supports. Once search is continued and stopped in accordance with Take Two or later, decisions are in an equally overwhelming majority made in accordance with Take Two. The two building blocks are therefore not independent but it is stopping that drives deciding. Stopping behavior is therefore a good predictor of the decision rule.

The third hypothesis did not receive any support. The small correlation between discomfort with ambiguity score and use of TTB's stopping rule even points in the opposite direction, with discomfort with ambiguity being rather associated with extended search beyond the first discriminating cue in the low-redundancy environment. One has to take into account, however, that there are at least two sources of ambiguity. One source is indeed the negative cue intercorrelations, experienced in the form of frequent contradictions between cues. The other source is ambiguity in terms of which alternative to choose. In the low-redundancy environment, a compensatory strategy achieves a much higher accuracy than TTB, so the price of extending search and facing the conflicting information is well worth paying, because one can find out how to handle this information successfully. Thus, using this information can actually decrease ambiguity about which alternative is correct. One factor that made the reduction of uncertainty about which alternative to choose possible is therefore the provision of outcome feedback in the learning phase.

This leads to a yet unanswered question: How did participants adapt their inference process? Participants in the high-redundancy environment could have experienced in the learning phase that decisions based on the first discriminating cue were just as accurate as decisions that take more information into account, such that a trade-off between accuracy and frugality was not necessary. Similarly, participants in the low-redundancy condition could have learned that integrating information is a much more accurate inference strategy than just relying on the first discriminating cue. Thus, adaptive strategy selection could be explained by learning the strategies' accuracies via outcome feedback. Participants might have learned which strategy performs better without even noticing the degree of information redundancy, that is, the degree to which the cues were correlated with each other. Alternatively, participants could have neglected the strategies' accuracies and simply learned that the cues were highly correlated with each other in the redundant environment, such that using a compensatory strategy would not lead to a different decision compared to a non-compensatory strategy. Since outcome feedback was given in the whole experiment it is impossible to decide how participants learned to select strategies adaptively. Thus, it is an open question whether people can select strategies that are adapted to the structure of decision environments without having had the opportunity to learn about the differential accuracy of strategies.

One hint that extensive learning with feedback is not always necessary is the fact that adaptive strategy use could already be observed in the first block of the decision phase, at a point when participants did not yet have experience about strategies' payoffs under conditions of information costs. They seem immediately to have taken costs into consideration and applied strategies that they anticipated would be well suited for the new situation. Admittedly, information costs are very salient and the benefits of using a frugal strategy are easy to see. This is in line with Payne et al. (1993) who suggest that information about strategies' costs might be more readily available than information about strategies' accuracies. Will people also be able to pick up the *correlational* structure of the environment, which mainly affects strategy accuracy, and respond adaptively to it when outcome feedback is not provided?

Experiment 2

People have expectations about which strategies are adaptive under various conditions: Chu and Spires (2003) asked people to indicate which of several proposed strategies they would use when, for example, there is limited time, or when the decision is very important. Although these situational descriptions are very general and do not include the aspect of information redundancy, the results indicate that people have a clear notion of when to select different strategies. Along these lines, research in the area of multiple cue probability learning indicates that "task feedback" can substantially influence people's inferences (Balzer, Sulsky, Hammer & Sumner, 1992; Balzer et al., 1994). Task feedback refers to information about the structure of the environment, for instance, about cue validities and the correlations between cues. It has been shown that task feedback increases performance above the level achieved in a no feedback condition, and additional information about the participants' accuracy did *not* further increase performance (Balzer et al., 1992).

Adaptive strategy use based solely on task information requires that people adequately perceive the correlational structure of the environment. Research on covariation assessment suggests that people sometimes have problems performing this task, especially in the case of binary events (see Alloy & Tabachnik, 1984, for a review). However, in many of these studies the degree of covariation had to be assessed, whereas for the inference task used in my studies, it was sufficient to detect whether and in which direction cues were correlated with each other. This seems to be a manageable task (e.g., Knowles, Hammond, Stewart & Summers, 1972). Moreover, under certain circumstances the covariation between binary events is judged accurately, for instance, when information is presented in summary tables (Ward & Jenkins, 1965). Also in Experiment 1 the detection of information redundancy might have been simplified because the cue values of all cues could be observed simultaneously for both alternatives. Thus, the frequent occurrence of contradictions between successive cues (i.e., cues supporting different alternatives) could be used as a shortcut to identify the low-redundancy environment, whereas frequent accordance between cues (i.e., cues supporting the same alternative) is indicative of the high-redundancy environment. One precondition for detecting these characteristics is that a large amount of information is looked up in the learning phase, which many participants in Experiment 1 did. Thus, it appears possible that even without outcome feedback people might be able to select inference strategies adaptively. Hypotheses 1 and 2 will therefore be tested again under a condition of learning without outcome feedback.

Hypothesis 3 will also be tested again. If the explanation suggested earlier for its failure to receive support in Experiment 1 is true, then the removal of feedback will also remove the possibility to decrease ambiguity about which option is correct through extended search and subsequent integration of information. Such a learning phase affords acquisition of knowledge about the correlational structure of the decision environment, but not about the differential accuracy of strategies. With this knowledge, only the ambiguity arising from successive cues contradicting each other can be countered, and it can be done by stopping search as soon as one discriminating cue has been found. The proposed association between ambiguity aversion and obeying TTB's stopping rule in the low-redundancy environment should therefore now be found.

Method

Participants

Forty participants (26 female, 14 male; average age 26 years), mostly students from various departments of the Free University of Berlin, took part in the experiment. For their participation, they received performance-contingent payment of on average 10.42 €, ranging from 2.50 to 17.97 €. The computerized task lasted approximately 1 hour. Afterwards, they filled out the German version of the Need for Cognitive Closure questionnaire (Deutsch et al., 2003), for which they received an additional 2.00 €.

Procedure

The task for the participants and the decision environment were identical to Experiment 1. The only aspect in which Experiment 2 differed from Experiment 1 was the learning phase. In the learning phase, no outcome feedback was provided. Participants were neither told whether their decision had been right or wrong nor provided with information about how much they had earned up to that point. This means that participants had no opportunity to learn the performance of the strategies. Outcome feedback was only introduced in the decision phase, which was identical to the decision phase of Experiment 1.

Results

Performance of participants

An analysis of variance with payoff as dependent variable, phase (learning vs. decision phase) as a within-subject factor, and environment as a between-subjects factor (high vs. low redundancy) was conducted. The average payoff of 10.69 € in the learning phase was higher than the average payoff of 2.17 € in the decision phase (see Table 2.5), $F(1, 38) = 441.31$, $p = .001$, $\eta^2 = 0.92$, which is easily explained by the fact that no search costs had to be paid in the learning phase. Overall, participants in the low-redundancy environment received on average higher payoffs (13.96 €) than participants in the high-redundancy environment (11.76 €), $F(1, 38) = 5.50$, $p = .024$, $\eta^2 = 0.13$. However, this effect only holds for the learning phase due to an interaction effect: In the learning phase participants in the low-redundancy environment received, with 11.66 €, on average higher payoffs than participants in the high-redundancy environment, with 9.72 €. This did not hold for the decision phase (2.30 vs. 2.04 €), $F(1,38) = 4.34$, $p = .044$, $\eta^2 = 0.10$. This can again be explained by higher accuracy of compensatory strategies in the low-redundancy environment than any strategy in the high-redundancy environment, with accuracy being the only determinant of the payoffs in the learning phase. When in the decision phase participants' performance became a function of both how accurate their inferences were and how frugal they were in their information search, the payoff differences between conditions disappeared.

Participants in both conditions achieved a similar accuracy level in the decision phase with 73% correct inferences in the low-redundancy environment and 72% correct inferences in the high-redundancy environment, $t(32) = 0.17$, $p = .433$ (one-tailed), $d = 0.05$. In both environment conditions, participants searched for the same amount of information on average, which was 2.3 cues. These results do not provide support for the adaptive strategy selection hypothesis. However, to test Hypotheses 1 and 2, it is necessary to analyze how well the strategies predict participants' decisions.

Table 2.5: Accuracy, frugality, and payoff participants achieved in the high- and low-redundancy conditions of Experiment 2, reported separately for learning and decision phase.

	High information redundancy		Low information redundancy	
	Learning phase	Decision phase	Learning phase	Decision phase
Mean percentage correct (*SD*)	75% (*2.9*)	72% (*5.7*)	80% (*7.3*)	73% (*9.4*)
Mean number of cues looked up (*SD*)	4.4 (*1.6*)	2.3 (*0.7*)	4.3 (*1.5*)	2.3 (*1.0*)
Mean payoff (*SD*)	9.72 € (*1.1*)	2.04 € (*1.6*)	11.66 € (*2.8*)	2.30 € (*1.9*)

Testing the strategies

First, it is tested how well TTB and the compensatory strategies predict the outcomes of participants' inferences, and second, how well they predict participants' information search.

In the learning phase, TTB predicted on average 92% of the inferences in the high-redundancy condition compared to 73% in the low-redundancy condition. In comparison, NB predicted 96 and 84% of the inferences in the high- and low-redundancy conditions, respectively. Thus, the compensatory strategy predicted the outcomes of more inferences in both environments.

In the decision phase, TTB predicted 88% of the inferences in the high-redundancy condition compared to 75% in the low-redundancy condition. NB predicted on average 88 and 76% of the inferences in the high- and low-redundancy conditions, respectively. Now, TTB and NB predicted a very similar proportion of decisions in both environments. Thus, no differences between strategies are observed in predicting participants' decision outcomes in the decision phase. This is to a large extent due to the high overlap in strategies' predictions, which makes these numbers difficult to interpret. Therefore, more importantly, differences in the search processes were examined to infer differential strategy selection.

I computed the order in which cues were looked up, and at which point search stopped. Whereas in the learning phase search order as well as frugal stopping should be of minor concern, in the decision phase frugality becomes crucial for increasing payoffs: Cues should be searched in the order of their validities, and search should be limited through adaptive stopping.

How did participants search for information? In the high-redundancy condition participants searched in the order of validity in 54% of the inferences in the learning phase compared to 75% of the inferences in the decision phase, $t(19) = 2.42$, $p = .013$ (one-tailed), $d = 0.54$. In the low-redundancy condition participants searched in the order of validity in 58% of the inferences in the learning phase compared to 69% of the inferences in the decision phase, $t(19) = 1.52$, $p = .073$ (one-tailed), $d = 0.34$. This small difference can be explained by the already high level of search according to validity in the learning phase. One reason why participants searched according to validity even in the learning phase might be that having a list of several cue values ordered by the cues' validities can simplify information processing. Nevertheless, in general, participants more frequently searched cues in the order of their validity when costs were introduced. Since the search order does not discriminate between strategies, the focus is on when search stops.

First, the proportion of inferences in which information search stopped before any discriminating cue was found (i.e., participants guessed) was determined. In the learning phase, search stopped before a discriminating cue was found in less than 1% of all inferences, while in the decision phase the proportion of guesses increased to 14%.

To test whether participants selected TTB for their inference or a compensatory strategy (e.g., NB or Take Two), I determined for all inferences in which participants did not stop their search before they found a discriminating cue whether search stopped when one discriminating cue was found (as predicted by TTB) or whether search continued (as predicted by the compensatory strategies). Consistent with Hypothesis 1, adherence to TTB's stopping rule increased substantially in both conditions when search costs were introduced: In the learning phase of the high-redundancy condition, participants stopped their search in 29% of the non-guessing inferences when one discriminating cue was found compared to 63% in the decision phase, $t(19) = 4.75$, $p = .001$ (one-tailed), $d = 1.06$ (see also Figure 2.3). In the learning phase of the low-redundancy condition, participants stopped their search in accordance with TTB in 19% of the non-guessing inferences compared to 42% in the decision phase, $t(19) = 3.15$, $p = .003$ (one-tailed), $d = 0.70$.

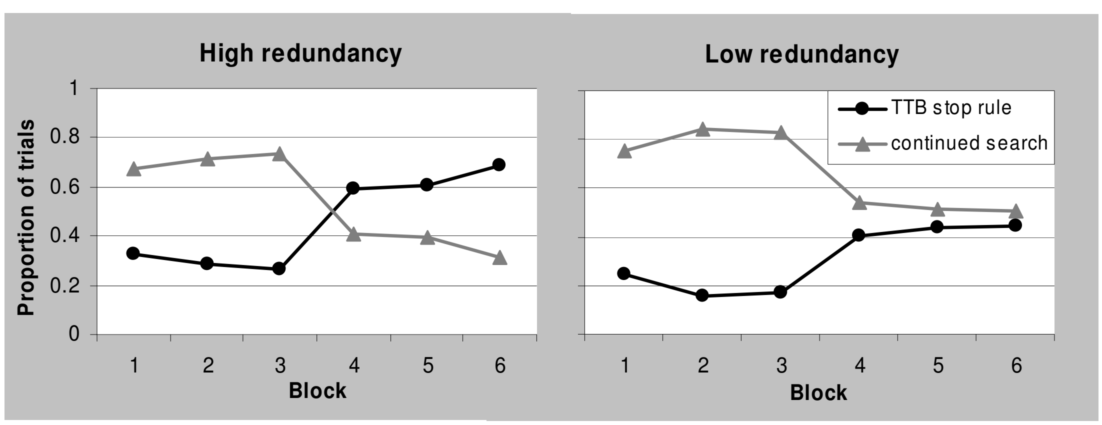

Figure 2.3: Proportion of non-guessing trials in which search stopped in accordance with TTB (i.e., when one discriminating cue was found) compared to the complement proportion of instances in which search continued beyond a first discriminating cue, across the six blocks of trials in Experiment 2.

The comparison between the decision phases of the high- and low-redundancy conditions in terms of when participants stopped their information search provides support for Hypothesis 2. Consistent with the hypothesis, the proportion of inferences in which search stopped when one discriminating cue was found (as predicted by TTB) was, with 63%, on average higher in the high-redundancy condition compared to the low-redundancy condition (with 42%), $t(38) = 1.96$, $p = .029$ (one-tailed), $d = 0.62$. This effect can even be found in the first block of the decision phase, block 4, with TTB-consistent stopping in 59% of trials in the high-redundancy condition compared to 41% in the low-redundancy condition. Nonetheless, in the high-redundancy environment the proportion of TTB-consistent stopping still increased across the decision phase, with on average 61% in block 5 compared to 68% in block 6, $t(19) = 2.18$, $p = .021$ (one-tailed), $d = 0.29$. Figure 2.3 shows the proportion of non-guessing trials in which search stopped in accordance with TTB's stopping rule, and the complement proportion of trials in which search continued beyond a first discriminating cue across the six blocks of trials for both redundancy conditions.

Similar to Experiment 1, in the high-redundancy condition search predominantly stopped after one discriminating cue was found, as predicted by TTB. In the low-redundancy condition search mostly continued, as predicted by compensatory strategies. Which of the two compensatory strategies predicts the information search better, Take Two or NB?

In the decision phase of the high-redundancy condition search stopped in accordance with Take Two in 31% of all non-guessing inferences compared to 27% in accordance with NB, $t(19) = 2.12$, $p = .024$ (one-tailed), $d = 0.48$. In the decision phase of the low-redundancy condition, Take Two predicted stopping in 27%, and NB in 24% of the non-guessing cases, $t(19) = 1.24$, $p = .115$ (one-tailed), $d = 0.43$. Take Two is thus better than or at least equally good as NB in predicting when participants stopped information search. Contrary to the results of Experiment 1, Take Two predicts the stopping behavior equally well in the low-redundancy condition (with 27%) and in the high-redundancy condition (31%), $t(37) = -0.52$, $p = .696$ (one-tailed), $d = -0.17$. As reported above, search continued beyond a first discriminating cue on the majority of inferences in the decision phase of the low-redundancy environment. The relatively low overall fit of Take Two (as well as NB) in the low-redundancy condition thus means that half of the cases in which search was continued beyond the first discriminating cue were *not* predicted by Take Two or NB, implying that participants often exhibited stopping behavior that was not captured by any of the suggested strategies. Exhaustive search does not explain the relatively high proportion of stopping that is inconsistent with TTB, Take Two, and NB: Participants looked up all cues in only 3% of the non-guessing inferences in the high-redundancy condition, and in 4% in the low-redundancy condition.

Is information search pattern a valid indictor of the strategy being used?

The results of Experiment 1 showed that stopping is a valid indicator of the decision rule that is applied. Does this also hold for Experiment 2? In the overwhelming majority of inferences

(99%) in the decision phase in which search was stopped when one discriminating cue had been found, decisions were in accordance with this first cue. I also selected all cases in the decision phase in which Take Two could be applied, that is, cases in which search continued until the point predicted by Take Two, or beyond. In the high-redundancy environment, decisions were made in accordance with Take Two in 98% of these cases. In the low-redundancy environment, where TTB and compensatory strategies more often make opposing predictions, the respective percentage was 96%. Even when only focusing on the subset of critical profiles, for which TTB and Take Two *differ* in their predictions (i.e. after one discriminating cue, two more are found that point in the other direction), results look similar. Participants decided in favor of the alternative to which the two cues pointed in 94% of these cases. These results demonstrate that also in Experiment 2, stopping is a valid indicator of the decision rule, and they support the notion that stopping drives deciding.

Information search and discomfort with ambiguity

As participants initially did not have the chance to learn from outcome feedback in the learning phase, the prediction was that discomfort with ambiguity would now be associated with higher levels of TTB use in the decision phase in the low-redundancy environment, as early stopping was the only possibility participants had to decrease ambiguity (Hypothesis 3). Indeed, a substantial positive correlation between TTB use and discomfort with ambiguity score on the Need for Cognitive Closure Scale was found for participants in the low-redundancy environment, r = .52, p = .009 (one-tailed). In the high-redundancy environment, the correlation was small, r = .29, p = .106 (one-tailed).

Discussion

In Experiment 2 people again adapted their inference process to both information costs and the information redundancy of the environment. In the decision phase, adaptive strategy selection in response to the degree of redundancy encountered in the environment was observed from the first block of the decision phase onward. Participants mostly searched cues in the order of their validities in both conditions. In the high-redundancy environment, after the introduction of information costs, one-reason stopping and deciding was observed in the majority of trials. In the low-redundancy environment, participants mostly continued search beyond a first discriminating cue. The effects, however, were weaker than in Experiment 1. Especially in the high-redundancy environment, stopping in accordance with TTB's prediction was observed less often in the decision phase compared to Experiment 1. Note that for the participants of Experiment 2 not only were information costs introduced after the decision phase, they also only then received outcome feedback. This could have led some participants to suspect that the environment might also have changed. With this assumption, increased exploration and neglect of TTB's stopping rule becomes reasonable, at least at the beginning of the decision phase. Consistently, the fit of TTB's stopping rule increased over the course of the decision phase, giving support to the notion that some participants only slowly started to rely solely on the first discriminating cue, although they could have already found out in the learning phase that looking up additional information rarely revealed “new” information in the sense of counterevidence against the first cue.

Among the cases of continued search, Take Two was again best in predicting when information finally stopped. However, many of the cases of continued search – about half in the low redundancy environment – were *not* captured by this heuristic, nor by the more complex strategy NB that takes cue weights into account. To understand this result, one needs to consider that the class of compensatory strategies, contrary to the class of non-compensatory strategies, is rather large. Once one-reason stopping is abandoned, there are many possible stopping rules and even more possible ways of deriving a decision from the encountered information. Besides the suggested stopping rules, people might sometimes try to stop earlier, by, for instance, looking up one more cue after the first discriminating cue and stopping search if it does not discriminate. Another alternative might be rules that stop after a fixed number of cues, always searching, for instance, for the three most valid cues. These examples give an impression of the variety of possible stopping rules that could be used for compensatory decision making. Thus, it is not very surprising that Take Two’s stopping rule obtains a relatively low fit. This also explains the replicated finding that TTB predicts decision processes better than Take Two and NB, even in the low-redundancy environment.

In sum, Experiment 2 demonstrates that people select strategies adaptively even when no outcome feedback is provided. This shows that the mere experience of different degrees of redundancy in a decision environment can trigger the selection of strategies that are well suited to deal with this kind of structure. Provision of outcome feedback, which enables learning of strategies’ accuracies, is thus not a necessary precondition. However, since the effects observed in Experiment 2 were smaller compared to Experiment 1, outcome feedback enhances adaptivity.

With regard to the third hypothesis about individual differences in the use of TTB, now indeed TTB-consistent stopping behavior was positively associated with expressed “discomfort with ambiguity” in the low-redundancy environment – an association that is, as expected, missing in the high-redundancy environment. This finding is in line with the explanation that lack of outcome feedback in the learning phase only left one way to decrease ambiguity: that of truncating search as soon as one discriminating cue had been found. TTB’s one-reason stopping rule thus spares ambiguity-averse participants encounters with conflicting information that would cause an ambiguity that this time could not be resolved through careful monitoring of outcome feedback.

General discussion

The first goal of this chapter was to test the performance of different strategies in situations of high versus low information redundancy. The second goal was to find out whether people adaptively select strategies in response to information redundancy in the decision environment. Thirdly, a potential correlate of individual differences in strategy use, discomfort with ambiguity, was measured to find out whether strategy use might also reflect personal likes and dislikes. Finally, adaptivity was studied under two conditions, with and without provision of outcome feedback, to test whether learning about strategies' differential accuracy was a necessary condition for observing adaptive strategy selection. What were the main findings of the three studies that addressed these questions, and what can be concluded from them?

The Simulation Study demonstrates that the most frugal model of the competition, the one-reason decision-making heuristic TTB, matches the accuracy of compensatory strategies even under adverse conditions such as low redundancy (as long as validity varies widely), or low cue validity dispersion (as long as redundancy is high). In cross-validation, TTB's performance proved to be very robust in these environments. The good relative performance in comparison to compensatory strategies in high-redundancy environments is not surprising, as under high redundancy, strategies will very frequently lead to the same predictions. The more astonishing finding is that TTB can avoid losses in accuracy relative to compensatory strategies even when redundancy is low if cue validities are widely dispersed. By relying on the most valid piece of information, TTB can obviously hold up to the accuracy of compensatory strategies that, when validities vary widely, will often not find enough counterevidence to overrule a highly valid cue encountered first. Thus, the benefit of information integration is diminished when validities vary widely. This finding is related to the analytical result that the accuracy of TTB is equivalent to that of a linear model with a non-compensatory set of weights (decaying in the same order as TTB's cue hierarchy; Martignon & Hoffrage, 1999). The more widely dispersed validities are, the closer cue weights based on these validities will come to a non-compensatory structure. This can explain why TTB's accuracy is very similar to that of compensatory strategies even under low redundancy as long as cue validities differ widely. Only when redundancy is low and at the same time cue validities are similar to each other does TTB's accuracy fall behind that of compensatory strategies, and losses in accuracy from calibration to validation increase.

On the other hand, compensatory strategies cannot generally be equated with lower robustness compared to one-reason decision making. Naïve Bayes, the most complex model in the competition, shows no signs of overfitting, suffering no more loss in accuracy when generalizing to new data than the other strategies. A particularly simple and quite frugal compensatory strategy does not fall far behind Naïve Bayes in terms of accuracy: Take Two. It stops search as soon as two cues are found that point to one alternative – thus also providing acceptable frugality – and decides accordingly. It achieves high accuracy especially under conditions in which TTB's performance breaks down: low redundancy and low validity dispersion. Hence, Take Two represents a very simple yet still quite accurate and frugal compensatory strategy that is as explicitly defined as TTB in its assumed decision process, that is, its search, stopping, and decision rules. The Simulation Study has thus identified two simple heuristics, likely to be in reach of people's cognitive capacities, that perform well under different degrees of information redundancy: TTB under high, and Take Two under low information redundancy.

How did participants respond to different degrees of redundancy in the decision environment? In both experimental studies, adaptive strategy selection was observed. People cut down their information search when costs for cues were introduced. More importantly, strategy use differed systematically between the two redundancy conditions, with a high percentage of TTB-consistent search patterns when redundancy was high, and a high percentage of search patterns consistent with compensatory decision making when redundancy was low. It is indeed Take Two (and not NB) that predicts stopping correctly in most of the decisions where search continued beyond the first discriminating cue. However, even in the conditions in which compensatory decision making yields higher payoffs than one-reason decision making, and TTB use is consequently relatively low, Take Two does not achieve as high a fit as TTB. Thus, although most participants did not adhere to TTB, there is no single compensatory strategy, as explicitly specified as TTB in its search, stopping, and decision rules, that predicts inference processes as well as TTB. Due to the high overlap between the predictions of different decision strategies – a characteristic that is inevitably associated with high information redundancy – strategies were mainly identified through the pattern of information search. Particularly stopping behavior proved to be a highly valid predictor of the decision outcome. When search was stopped when one discriminating cue was encountered, participants for the most part relied on this piece of information. When instead more information was acquired, participants also integrated this information to arrive at a decision.

One property of TTB is that it avoids encountering conflicting information. When search is stopped after the first discriminating cue has been found, no conflict between cues can arise. For some people, this characteristic might make TTB an attractive heuristic to use when cues are likely to contradict each other. In particular, TTB should be especially attractive for people who find ambiguity aversive – even when compensatory strategies would achieve higher accuracy, as is the case in the low-redundancy environment of the present studies. When participants had the opportunity to learn about strategy accuracy through outcome feedback, I did not find evidence for the hypothesized association between TTB use and self-reported discomfort with ambiguity on the Need for Cognitive Closure Scale (Kruglanski et al., 1993). In hindsight, this result does not appear surprising: Through learning that compensatory strategies are much more accurate than TTB in the low-redundancy environment, participants could actually also decrease ambiguity – namely ambiguity about which alternative to choose. Thus, they might have willingly faced information in the form of cues that often contradicted each other because they knew what to do with this information and how to integrate it to derive a decision that would very likely be correct. In contrast, when in Experiment 2 outcome feedback was not provided, TTB use was positively associated with self-reported discomfort with ambiguity in the low-redundancy environment. Seemingly, when not being able to find out what strategy works best in terms of accuracy, and when cues often contradict each other, TTB becomes an attractive strategy for conflict-averse people.

This result suggests that individual differences might represent a fruitful area for further research. It could help to address the problem of often large individual differences in strategy use even within the same environmental condition. In the preference domain, individual differences in choice processes have given rise to very stimulating research, with possible implications for applied questions in the area of clinical psychology as well as consumer psychology (e.g., Schwartz et al., 2002). Research on fast and frugal heuristics for inferences has so far focused on the explanatory potential of the structure of the decision environment. Compared to the growing number of demonstrations of the effect of environmental characteristics on strategy use, little attention has been paid to differences among decision makers. Bröder (2003) provides an exception, demonstrating a moderating effect of intelligence on adaptive strategy selection. The fact that in Experiment 2 participants who reported higher levels of discomfort with ambiguity tended to use TTB more frequently suggests that also differences in personal likes and dislikes might influence strategy selection. Taking into consideration dimensions of individual differences in studies on inferences in addition to environmental variables might help us achieve a more comprehensive picture of decision makers.

The fourth main question was whether adaptivity depends on the provision of outcome feedback, or if it would be observed even when participants did not have the opportunity to learn about the differential accuracy of strategies. Experiment 2 provides evidence for adaptive strategy selection without the previous opportunity to learn from outcome feedback, thus suggesting the possibility of more "intuitive" forms of adaptivity. A helpful distinction has been suggested by Payne et al. (1993), between strategy selection depending on the *anticipated* accuracy and costs of strategy application (top-down approach) and selection depending on the *experienced* accuracy and costs of the strategies, based on the information encountered during the decision process (bottom-up approach). Rieskamp and Otto (2004) have shown that people can learn through outcome feedback to apply strategies that achieve high accuracy in a given environment. However, it cannot be assumed that people regularly get the chance to go through a long series of learning trials with feedback to find out about strategies' performances. Many decisions lack the opportunity to learn slowly from outcome feedback, or they provide only late feedback, like, for instance, in the case of decisions about pension plans. Then, people have to anticipate which strategies are likely to be successful. Certain environmental structures could themselves serve as cues for selecting particular strategies, providing information about the likely performance of those strategies. Degree of redundancy of an environment, which might be perceived via a shortcut, namely the existence of frequent cases of accordance, or contradictions, respectively, between cues, seems to be treated as a hint to select particular strategies.

Strategy selection does not presuppose a conscious deliberation process. Payne et al. (1993) suggested that strategy selection can be a conscious choice but more often reflects a learned contingency between certain task characteristics and strategies' accuracy and costs. It is also not ruled out that some of these contingencies might even be hardwired and the product of evolutionary processes. But in any case, the assumption that environmental characteristics per se suggest the selection of certain heuristics allows us to circumvent the dilemma that although fast and frugal heuristics are very easy to apply, the problem of how they are selected either remains unsolved, or is addressed by reinforcement learning theories that predict a gradual and relatively slow learning process based on feedback (Rieskamp & Otto, 2004). Many authors (Feeney, 2000; Luce, 2000; Morton, 2000; Newstead, 2000; Wallin & Gärdenfors, 2000) have criticized work on fast and frugal heuristics (as summarized in Todd & Gigerenzer, 2000) for not specifying how heuristics are selected from the "adaptive toolbox" in response to environmental characteristics. Although to provide a satisfactory answer to the selection problem, clearly more has to be learned about the origins of people's perceptions of strategy-environment contingencies, my results suggest that adaptive strategy selection can to a large part be explained without invoking a slow reinforcement learning processes. Information costs and correlational structures influence people's expectations about which strategy will perform well in a particular environment and consequently affect strategy selection even when there is no opportunity to learn from outcome feedback how well, exactly, a particular strategy performs. Nevertheless, it has to be noted that availability of outcome feedback resulted in more adaptive strategy selection on average. Thus, expectations about strategies' performances are updated when outcome feedback is available, which yet advances people's adaptive inference processes.

However, the conclusions from the experimental results have to be restricted. First, the effect of information redundancy was strongest when search costs were introduced. In contrast, when information was for free (i.e., in the learning phase), participants often searched extensively for information, contrary to the predictions of simple strategies. Although this behavior is still adaptive, since it did not reduce participants' payoffs, extensive search did not increase participants' payoff either and could have been omitted. It is an open question whether participants would restrict their information search after extensive learning even when information search does not involve any costs.

Second, in both experiments participants made inferences in artificially constructed instead of natural environments. As explained earlier, in the domain of inferences cues usually – that is, in natural environments – correlate positively with each other due to the constraint that useful cues are positively correlated with the criterion. Thus, especially the low-redundancy condition of Experiment 1 and 2, with slightly negative correlations between cues, might have represented an unnatural inference situation. Because of this, it could have been possible that people have difficulties reacting to low information redundancy in adaptive ways. However, inference situations in natural environments do vary to the degree of redundancy. Thus, although the low-redundancy condition might have represented a *relatively* unfamiliar environment, information redundancy is a critical dimension on which inference problems differ and in fact, adaptive responses to information redundancy were observed.

Finally, the results of my experiments show that people respond adaptively to information redundancy even when no outcome feedback is provided. But it is an open question whether people are also able to adapt their inference process to other environmental characteristics when no outcome feedback is provided. Environmental characteristics like the "compensatoriness" of cue weights or the scarceness of information (Martignon & Hoffrage, 1999) can influence strategies' accuracies. Bröder (2000, Experiment 2), for instance, investigated the effect of dispersion of validities on strategy selection, without conclusive results. However, the inference situation in his experiment did not involve any information search, as inferences were made from givens. It is therefore an open question whether dispersion of cue validities influences inference processes under ecologically more valid circumstances that require that information has to be searched for.

Conclusion

The results of the present studies suggest that people select strategies adaptively in response to environmental characteristics. Their inference processes differ considerably depending on whether they encounter inference situations with high or low information redundancy. Feedback about the accuracy of their decisions supports the selection of adequate strategies. Yet even without outcome feedback, people intuitively select different strategies under different circumstances. Accumulation of results such as these might substantiate arguments against the criticism that the adaptive toolbox approach to human decision making has merely replaced the problem of applying a complicated yet general decision making strategy by the problem of selecting an adequate heuristic from a whole box of simple rules (Newell et al., 2003). The present results demonstrate that adaptive strategy selection does not necessarily require individuals to slowly learn from outcome feedback how accurate strategies are in a particular environment, suggesting that the selection process can be driven by people's perceptions of various structural aspects of environments. It should be a fruitful and interesting enterprise to explore the origins of such contingency knowledge.