Memory & Cognition
2007, 35 (7), 1801-1813

The influence of information redundancy on probabilistic inferences

Anja Dieckmann and Jörg Rieskamp
Max Planck Institute for Human Development, Berlin, Germany

Information redundancy affects the accuracy of inference strategies. A simulation study illustrates that under high-information redundancy simple heuristics that rely on only the most important information are as accurate as strategies that integrate all available information, whereas under low redundancy integrating information becomes advantageous. Assuming that people exercise adaptive strategy selection, it is predicted that their inferences will more often be captured by simple heuristics that focus on part of the available information in situations of high-information redundancy, especially when information search is costly. This prediction is confirmed in two experiments. The participants' task was to repeatedly infer which of two alternatives, described by several cues, had a higher criterion value. In the first experiment, simple heuristics predicted the inference process better under high-information redundancy than under low-information redundancy. In the second experiment, this result could be generalized to an inference situation in which participants had no prior opportunity to learn about the strategies' accuracies through outcome feedback. The results demonstrate that people are able to respond adaptively to different decision environments under various learning opportunities.

Decisions are often made under uncertainty. People cannot perfectly predict whether their chosen route to work will avoid the morning traffic jam or whether their favorite basketball team will win an upcoming game. Thus, they have to rely on cues—pieces of information that are imperfectly correlated with the criterion to be predicted. Such cues are sometimes highly correlated with each other and speak for the same prediction, whereas at other times they contradict each other and suggest incompatible predictions. This variance in information redundancy in decision environments is the focus of this article.

How people's inferences are influenced by information redundancy has been addressed in the neo-Brunswikian "social judgment theory" research (Brehmer & Joyce, 1988; Cooksey, 1996; Doherty & Kurz, 1996). This is not a coincidence. One of Brunswik's (1952, 1955) core concepts is vicarious functioning. Broadly, this can be defined as "exchangeability of pathways relative to an end" (Brunswik, 1952, p. 17). More narrowly applied to the problem of probabilistic inference, vicarious functioning means that correlated cues can function as substitutes for each other. In social judgment theory research, experimenters have often applied the multiple cue probability learning (MCPL; Smedslund, 1955) paradigm. Typically in this line of research, people's judgments are described by fitting a regression model to the judgment and by comparing the model's resulting beta weights with the beta weights of the "ecological" regression model fitted to the correct criterion values (Brehmer, 1974). Information redundancy, defined as 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). However, Armelius and Armelius (1974; see also Schmitt & Dudycha, 1975) doubted that people take correlations between cues into account when making inferences. They showed that the beta weights of regression models fitted to participants' inferences matched cue–criterion correlations rather than the ecological beta weights incorporating correlations between cues.

We argue that this skepticism about people's ability to incorporate cue redundancy into judgments has derived from using mainly one model for describing people's judgments—namely, linear regression. We suggest that people's judgments may rely on inference strategies that cannot easily be captured by a regression model. For instance, under high-information redundancy people might select strategies for their judgments that ignore information, since these strategies will still perform well. When fitting a regression model to these judgments, the resulting beta weights might deviate from the beta weights of the ecological regression model. Thus, although a decision maker behaves adaptively in response to information redundancy by selecting appropriate inference strategies, this might not be reflected in the fitted regression model. Hammond (1996) admits that the predominant use of multiple regression might have prevented consideration of alternative models that are more accurate in describing human decision-making processes, including the possibility that people react to information redundancy by using

A. Dieckmann, anjadieckmann@gmail.com

specific inference strategies (see also Dhami, Hertwig, & Hoffrage, 2004).

In general, many authors have argued that people are equipped with a repertoire of decision strategies from which to select, and which one is chosen depends on the decision context (e.g., Einhorn, 1970, 1971; Gigerenzer, Todd, & the ABC Research Group, 1999; Ginossar & Trope, 1987; Payne, 1976; Payne, Bettman, & Johnson, 1988, 1993; Rapoport & Wallsten, 1972; Rieskamp, 2006; Rieskamp & Otto, 2006; Svenson, 1979). In the domain of preferential choices, Bettman, Johnson, Luce, and Payne (1993) provided support for the notion that people also select strategies adaptively in response to information redundancy. They showed that participants choosing between gambles searched only for a subset of the available information when they encountered a redundant environment with positively correlated attributes. Negatively correlated attributes, in contrast, gave rise to search patterns consistent with compensatory strategies that integrate more information. However, the converse pattern was found by Luce, Bettman, and Payne (1997) for emotionally difficult choices, such as job choice, where negatively correlated attributes induced participants to focus more on one attribute at a time, which is in line with the selection of simple, noncompensatory decision strategies (see also Luce, Payne, & Bettman, 1999). Lastly, Johnson, Meyer, and Ghose (1989), for choices between hypothetical apartments, did not find any evidence for the selection of different strategies in response to changes in interattribute correlation. Hence, results in the preference domain are rather mixed and it is not clear whether, and in which direction, information redundancy influences people's strategy selection. Additionally, preferences depend on people's subjective evaluations, and there is no consistent outside criterion that can be used to evaluate whether the selection of a certain strategy is adaptive. For inferences, such an outside criterion does exist, which allows us to strictly evaluate people's adaptivity.

Redundancy is also more typical of inferential rather than preferential decision making. Consider a preferential choice such as choosing among laptops. Here, the alternatives' attributes are subjectively evaluated as good or bad. Thus, these evaluations can be correlated in any direction without constraint, and negative correlations can be common, as, for instance, when a laptop's attractive large screen comes with an inconveniently high weight. In contrast, in an inference situation, such as predicting the winner of a basketball game, the cues describing the alternatives, such as, for instance, home advantage of one team, are predictors of the criterion to be judged. To be predictive, cues have to correlate positively with the criterion, which puts a lower limit on the intercue correlations. It is thus more likely that cues will be positively correlated with each other, representing a case of cue redundancy. Because information redundancy is a more common characteristic of inferences, people might be familiar with it and be more likely to respond to it adaptively. This expectation follows Brunswik's (e.g., 1955) idea of vicarious functioning as an essential principle of human achievement. In sum, we consider inferences an area in which adaptive strategy selection in response to information redundancy is likely to occur.

The following three studies pursued three goals. First, in a simulation study, we explored whether simple inference strategies can indeed be as accurate as more complex ones when information redundancy is high. Additionally we examined the amount of information various strategies require. From these results, we conclude that simple inference heuristics should predict people's inferences well under the condition of high-information redundancy. The second goal was to test this prediction with two experimental studies. These experiments also addressed our third aim: to explore whether adaptive strategy selection is restricted to situations with explicit outcome feedback.

Simple Heuristics for Probabilistic Inferences

Consider the problem of inferring which of two potential oil fields will yield more oil. For this inference, different tests could be carried out—chemical analyses, for instance, to determine the amount of organic matter in the bedrock (Mobil Oil, 1997). But which and how many tests should be considered, and how should test results be used to make an inference? Gigerenzer and Goldstein (1996) have suggested a simple heuristic called *Take the Best* (TTB) for inferring which of two alternatives, described by several binary cues, has a higher criterion value. TTB searches sequentially through cues in the order of their validity. The validity v can be defined as the conditional probability that a cue will lead to a correct inference—on the condition that it discriminates. A cue discriminates if the two alternatives have different cue values. Search is stopped as soon as one cue is found that discriminates between alternatives, and the alternative with the positive cue value is selected. If no cue discriminates TTB chooses randomly. This inference strategy is "noncompensatory" because a cue cannot be outweighed by any combination of less valid cues, in contrast to a "compensatory" strategy, which integrates cue values. Czerlinski, Gigerenzer, and Goldstein (1999) showed that TTB is very accurate in solving various inference problems.

Recent studies have explored under what conditions people select TTB (e.g., Bröder, 2000, 2003; Bröder & Schiffer, 2003b; Newell & Shanks, 2003; Newell, Weston, & Shanks, 2003; Rieskamp, 2006; Rieskamp & Hoffrage, 1999, in press; Rieskamp & Otto, 2006). All this work relies on the *adaptive strategy selection hypothesis*, defined as the assumption that strategies that perform well in a particular situation are most likely selected by the decision maker. For the purpose of this article we define a strategy's performance purely in monetary terms, that is, the payoff a strategy achieves when consistently applied. The payoff depends on search costs and monetary gains for accurate decisions. Thus, for simplification, the main source of a strategy's costs is assumed to depend on the strategy's required amount of information—its "frugality." Following this definition, according to the adaptive strategy selection hypothesis people should select the strategy from the strategy repertoire that produces the highest monetary payoff in a particular environment. This hypothesis simplifies the cost–benefit approach to strategy selection (Beach &

Mitchell, 1978; Christensen-Szalanski, 1978; Payne et al., 1988, 1993). The cost–benefit approach states that people trade off strategies' costs against their benefits when selecting a strategy. The strategies' costs can be related to the effort to process them or to the amount of information they require, and the strategies' benefits are related to their accuracy.

Past research has shown that simple heuristics indeed predict inferences well when the costs for applying compensatory strategies are high, for instance, when people have to pay for information acquisition (Bröder, 2000; Newell & Shanks, 2003, Newell et al., 2003; Rieskamp & Otto, 2006), make inferences under time pressure (Rieskamp & Hoffrage, 1999), or retrieve information from memory (Bröder & Schiffer, 2003b). More rarely, adaptive strategy selection has been demonstrated when strategies' accuracies were varied (Bröder, 2003; Rieskamp & Otto, 2006). In particular, the influence of information redundancy on strategies' accuracies, and, ultimately, on strategy selection, has garnered very little attention.

SIMULATION STUDY

The goal of the simulation study was to examine how information redundancy affects the accuracy and frugality of inference strategies. The inference task can be formalized as follows: An environment consists of N objects; each object i is characterized by a criterion value x_i. The task is to predict for all possible pair comparisons which object has the larger criterion value. Each object is described by a set of M binary cues. Each cue m can have a positive or nonpositive cue value c_m (i.e., 1 or 0). The cue value of the first alternative, A, minus the cue value of the second alternative, B, results in the cue difference d_m. Information redundancy occurs when the cues are correlated with each other. We describe three inference strategies that differ in computational complexity and information demands and show how their accuracy and frugality are affected by information redundancy.

Strategies for Probabilistic Inferences

The first strategy we consider is TTB, which does not integrate any information. The second strategy is *naive Bayes* (NB; see, e.g., Lee & Cummins, 2004), representing the class of compensatory strategies that integrate information of all available cues. Lee and Cummins (2004) argued that NB is the "rational" model for probabilistic inferences. In fact, Martignon and Laskey (1999) showed that NB outperforms TTB across a large number of real environments. Bröder (2000), Newell and Shanks (2003), and Newell et al. (2003) have used NB to determine which alternative was considered the correct inference in their experiments and provided outcome feedback 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 by adding the log odds for each cue (derived from the cue validities), multiplied by the cue difference:

$$\ln\left(\frac{p_k\left(x_A > x_B\right)}{1 - p_k\left(x_A > x_B\right)}\right) = \ln\left(\frac{v_1}{1 - v_1}\right)d_1 + \cdots + \ln\left(\frac{v_m}{1 - v_m}\right)d_m + \cdots + \ln\left(\frac{v_M}{1 - v_M}\right)d_M, \quad (1)$$

where k is a particular pair comparison and v_m is the validity of cue m (see also Lee & Cummins, 2004). If the predicted score is larger than 0, then the probability that the first alternative has a larger criterion value than 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 makes the simplifying assumption that cues are independent, thereby ignoring correlations between cues.

Due to NB's complexity, it is questionable whether people follow its predictions. Therefore, we also considered an alternative compensatory strategy that is simple to process, which we call *Take Two*. This strategy builds a bridge between TTB and NB: Like TTB, it searches for cues in the order of their validity. But unlike TTB, it only stops search when two cues that favor the same alternative have been found. This alternative is then selected regardless of whether, during search, another cue was encountered 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 founded on the idea that people might not be comfortable basing their decision on one single cue—and indeed people often do search beyond the first discriminating cue (e.g., Newell & Shanks, 2003)—but nevertheless they might want to limit their information search. Take Two satisfies these motivations.[1]

Testing the Strategies

The strategies' accuracies were mainly evaluated by their generalizability (Pitt & Myung, 2002), that is, their ability to make accurate inferences for a validation environment that was not used to estimate the strategies' parameters. Therefore, 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. Moreover, the strategies' levels of frugality were evaluated, defined as the average number of cues required to make an inference. Whereas TTB and Take Two define in which order to search for cues and when to stop, this is not clear for NB. At first glance, one would expect NB to require all available cues. But even for NB, limited information search is, in principle, possible. If, for instance, the five most valid of six cues favor one alternative, the sixth cannot change the preliminary decision and, therefore, does not need to be acquired. Thus, search can be limited by assuming that NB also searches for cues in the order of their validities

Table 1
Average Cue Validities and Average Correlation Between Cues in the Four Decision Environments of the Simulation Study

	High Redundancy		Low Redundancy	
	HVD	LVD	HVD	LVD
Validity				
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 (r) between cues	.51	.51	.01	.01

Note—HVD, high-validity dispersion; LVD, low-validity dispersion.

and stops when additional cues cannot change a preliminary decision based on the cues acquired so far. It is, of course, questionable whether this search process is psychologically plausible, since preliminary decisions have to be determined after each cue is acquired and compared to a hypothetical final decision. Nevertheless, we assumed this limited search for NB because it leads to a more demanding competition between the strategies in terms of frugality and enabled a more conservative test of TTB's apparent frugality advantage.

Method

Accuracy and frugality were tested in environments with either high or low-information redundancy. In addition, the dispersion of cue validities was varied: While the average cue validity was the same, validities varied widely in the high dispersion condition, and only moderately in the low dispersion condition. The dispersion of cue validities should also influence strategies' accuracy: Everything else being equal, relying on the most valid cue when it has a much higher validity compared to the remaining cues is more reasonable than when the cues have similar validities (Martignon & Hoffrage, 2002). The two factors, information redundancy and validity dispersion, were crossed, providing four conditions.

In more detail, 500 artificial environments were created randomly for each of the four conditions. Each environment consisted of 50 objects and 6 cues. Twenty-five positive (i.e., a cue value of 1) and 25 nonpositive (i.e., a cue value of 0) cue values were randomly assigned to each of the six cues, resulting in a discrimination rate of .51. To obtain the specified cue validities, as well as high and low average correlations between the cues (see Table 1), we applied a hill-climbing search process (described in the Appendix) that modified the environments repeatedly until the requirements were met. For the purpose of cross-validation half of the 50 objects of each of the 500 environments per condition were randomly selected to function as calibration sample to estimate the strategies' parameters, while the remaining 25 objects constituted the validation sample to test how well the strategies generalize.

Results

The strategies' accuracies and frugalities are shown in Table 2. Results are reported separately for the calibration and validation samples, for the low- and high-information redundancy, and for the low- and high-validity dispersion conditions. All strategies reached higher accuracy under low-information redundancy compared to high-information redundancy. More importantly, the strategies' accuracies in comparison to each other depended on the environment. When comparing accuracies for the crucial validation sample, TTB, the most frugal strategy in the competition, approximately matched NB in three environment conditions: Under high-information redundancy, regardless of the dispersion of the cue validities, and under low-information redundancy, if the dispersion of cue validities was high. Only when redundancy was low *and* cues had similar validities did TTB suffer a clear loss in accuracy, lagging eight percentage points behind NB. In this situation, Take Two matched NB's accuracy. In all other conditions as well, Take Two's performance approximated that of NB and TTB.

Robustness—that is, the strategies' losses in accuracy from calibration to validation sample—varied between one and eight percentage points. NB performed fairly robustly, with a maximum accuracy loss of four percentage points under the condition of low-information redundancy and low cue validity dispersion. TTB experienced similar accuracy losses, with the exception of the condition of low-information redundancy and low cue validity dispersion, where its accuracy dropped by eight percentage points. Take Two was the most robust strategy, with a maximum accuracy loss of two percentage points.

The number of cues acquired by the strategies to arrive at a decision was averaged across the calibration and validation samples, because the results for the two samples did not differ. Although limited information search was assumed for all strategies, making the competition more demanding, TTB clearly required the least information. With the assumption of limited information search, NB was slightly more frugal than Take Two, which can be attributed to the restricted number of cues (six). In condi-

Table 2
Mean Proportions of Correct Inferences in the Calibration and Validation Samples, and Average Numbers of Cues Looked Up (Frugality) in the Four Decision Environments

	Accuracy				Frugality	
	Calibration		Validation			
	M	*SD*	*M*	*SD*	*M*	*SD*
High Redundancy						
HVD						
TTB	.72	.03	.70	.04	2.3	0.2
Take Two	.69	.04	.68	.04	3.9	0.1
NB	.71	.03	.70	.04	3.4	0.3
LVD						
TTB	.71	.04	.68	.05	2.3	0.2
Take Two	.68	.04	.68	.05	4.0	0.1
NB	.70	.04	.68	.05	3.9	0.3
Low Redundancy						
HVD						
TTB	.85	.03	.82	.05	2.0	0.1
Take Two	.85	.04	.83	.05	4.2	0.1
NB	.87	.03	.84	.04	3.5	0.4
LVD						
TTB	.82	.03	.74	.05	2.0	0.1
Take Two	.84	.03	.82	.04	4.3	0.1
NB	.86	.03	.82	.05	4.0	0.5

Note—HVD, high-validity dispersion; LVD, low-validity dispersion; TTB, Take the Best; NB, naive Bayes.

tions with more cues, Take Two would become more frugal relative to NB, because regardless of the number of cues, it will stop search once two cues that support one alternative are found.

Discussion

Strategies' accuracies are affected by information redundancy. The prediction that TTB would perform accurately under high-information redundancy whereas the compensatory strategies would be more accurate under low redundancy generally holds. Surprisingly, as long as cue validities varied substantially, TTB was not outperformed by the compensatory strategies even under low redundancy. Obviously, a high dispersion of cue validities made it difficult for cues lower in the validity hierarchy to compensate for a cue of higher validity. Thus, TTB achieved a good cost–benefit ratio in most situations by being both accurate and very frugal. Only when information redundancy was low and cues had similar validities was TTB outperformed by the compensatory strategies. These results are important, as they demonstrate that calling NB the "rational" model (see Bergert & Nosofsky, 2007; Lee & Cummins, 2004) in comparison to TTB can be misleading. One could argue that it is more rational to apply TTB under high redundancy because it obtains the same accuracy as NB by using only half of the information NB requires.

EXPERIMENT 1

Experiment 1 tested the adaptive strategy selection hypothesis, which proposes that people should select the best-performing strategy for a decision environment. The simulation study showed differences in the strategies' frugalities, so that in a situation in which the performance of a strategy depends crucially on its information costs, people should limit their information search and select TTB more often. Additionally, the simulation study demonstrated that TTB is a very accurate strategy in situations of high-information redundancy. It can therefore be predicted that people should select TTB for making their inferences when encountering high-information redundancy. In contrast, in situations of low-information redundancy, compensatory strategies—represented by Take Two and NB—should be selected. This leads to two hypotheses:

Costs hypothesis: 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.

Redundancy hypothesis: 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.

To test these hypotheses we conducted an experiment in which participants first explored the environment in a learning phase without any information costs, which should enable them to select the most adaptive strategy for a subsequent decision phase, in which information was costly. To focus on the effect of cue redundancy on strategy selection, we did not let participants learn the cue validities over repeated trials with outcome feedback. Instead, we provided the respective validities directly, similar to in previous studies (e.g., Bröder, 2003; Newell et al., 2003; for a discussion of this issue, see Dieckmann & Todd, 2004).

Method

Participants. Forty participants, mostly students from various departments at the Free University of Berlin, took part in the experiment. They received performance-contingent payment for their participation of on average €10 (corresponding to ca. $10 at the time of the study), ranging from €2 to €19. The computerized task (Czienskowski, 2004), which was conducted in individual sessions, lasted approximately 1 h.

Procedure. Participants were asked to imagine they were geologists hired by an oil-drilling company. Their task was to decide at which of two potential drilling sites, labeled X and Y, more oil would be found based on various tests (such as chemical analysis). They could conduct up to six tests by successively clicking on corresponding icons on the computer screen. The test results were revealed for both drilling sites and remained visible until a decision was made (see Figure 1). Below each icon, the test's validity—labeled "success probability"—and direction were indicated (i.e., which of the two possible test outcomes pointed to larger oil deposits). The concept of cue validity was explained to participants.[2] The positions at which the tests were displayed on the screen were the same throughout all decisions for one participant but varied randomly across participants. The test results appeared on the lower part of the screen in the order in which the tests had been selected. After a decision had been made, outcome feedback was provided by either a green "correct" box or a red "wrong" box. Participants' payoffs were presented on the screen, expressed in an experimental currency called "petros," with ten thousand petros corresponding to €1. Participants earned €0.20 for a correct decision and paid €0.20 for a wrong decision.

The experimental design had two factors: phase (within subjects; learning phase vs. decision phase) and environment (between subjects; high redundancy vs. low redundancy). Whereas in the learning phase participants could look up information for free, in the decision phase they had to pay €0.03 for each test. The initial learning phase and the final decision phase consisted of 96 decisions each, containing three repetitions of blocks of 32 items. Within each block, items were randomly ordered (with the same order for all participants). The position of the correct drilling site for each item (left or right on the screen) also varied randomly.

More specifically, we constructed an experimental item set (different from the environments of the simulation study) with 32 items, each consisting of two alternatives described by six cues. Thus, each item could be represented by a cue difference vector, by subtracting the cue values of one alternative from the cue values of the other alternative. Cue difference values were assigned such that the validities of the cues were .83, .78, .72, .67, .61 and .56, and the discrimination rates had an equal value of .56 for each cue. For each cue, it was randomly determined in which of the 32 pair comparisons the cue would discriminate correctly, incorrectly, or not at all. To generate the high and low-information redundancy environments, cue difference values were switched between two randomly selected pair comparisons. If the switch increased the average intercue correlation (or decreased it, respectively, for the low-information redundancy condition), the change was accepted, if not, the change was undone.[3] This last step was repeated until convergence in the average intercue correlation was obtained. In the high-redundancy condition, the average correlation between the six cues was $r = .50$ (with a minimum correlation between two cues of $r = .35$) compared to $r = -.15$ in the low-redundancy condition (with a maximum correlation between two cues of $r = -.02$). Thus,

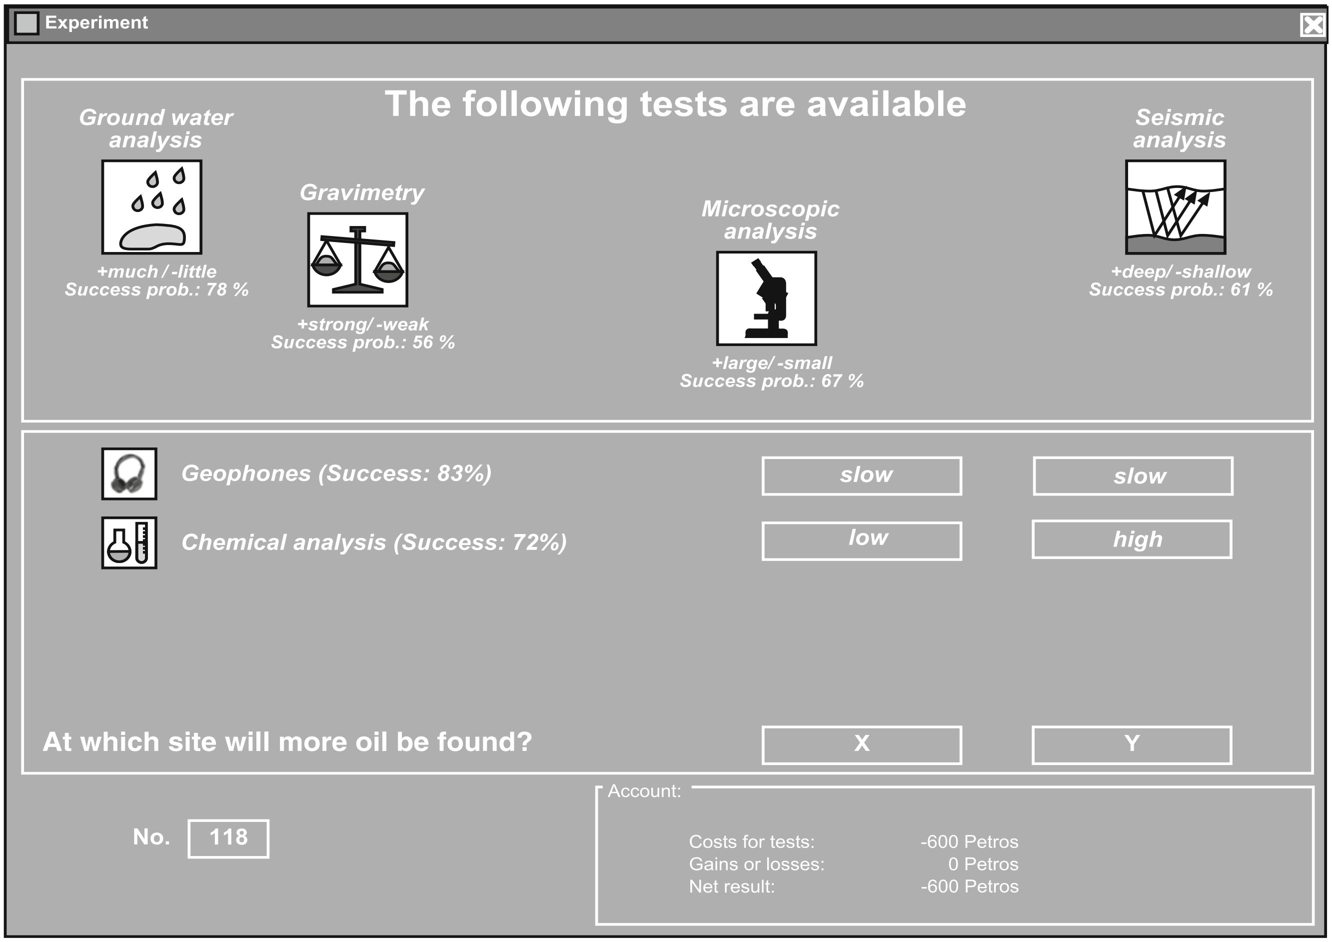

Figure 1. Screenshot of the computerized task participants faced in Experiments 1 and 2 (taken from the decision phase and translated from German).

under low redundancy the cues not only provided additional valid information, but also revealed pieces of information that were often in conflict with each other.

To confirm that our hypotheses were justified, we determined for both experimental environments how well the strategies would perform if applied consistently. NB and Take Two always predicted the same inference outcomes. In the high-redundancy environment, consistent with the results of the simulation study, TTB had the same accuracy as the two compensatory strategies, with 78% correct inferences. Thus, in the learning phase with no search costs all strategies reached the same monetary payoff (€10.80, added across all three blocks). In contrast, in the decision phase with search costs, TTB produced a higher payoff than NB (€4.77 compared to €1.62). In the low-redundancy environment, again in line with the simulation results, TTB achieved 66% correct inferences and was thus less accurate than NB and Take Two with 91% correct inferences. This leads to unequal payoffs already in the learning phase, where TTB produced a payoff of €6 compared to €15.60 for NB. In the decision phase, TTB led to a payoff of €1.41 compared to €4.44 for NB. Thus, the performance of TTB and the compensatory strategies was reversed in the decision phases of the low- and high-redundancy conditions. According to the redundancy hypothesis people should thus select TTB in the high-redundancy condition, and NB or Take Two in the low-redundancy condition.

Due to the redundancy between cues, the overlap in strategies' predictions was large in the high-redundancy environment, with on average 94% identical predictions of TTB and the compensatory strategies. The corresponding number for the low-redundancy environment was 63%. High overlap between different strategies' predictions is an immanent characteristic of high-redundancy environments. Therefore, we focused on process measures instead of inference outcomes. This approach has been demonstrated to be a valid alternative for identifying the inference strategies people select (see Bergert & Nosofsky, 2007; Bröder & Schiffer, 2003a; Payne et al., 1988).

Results

To test the costs and redundancy hypotheses, we determined first how well the different strategies predicted participants' inferences, and second how well they predicted the information search, that is, the order in which cues are acquired and when the search for information stops. Table 3 summarizes the results for the costs hypothesis, and Table 4 for the redundancy hypothesis.

To examine the costs hypothesis, we compared the proportion of participants' inferences predicted by the different strategies in the learning phase and in the decision phase. TTB predicted on average 79% of the inferences in the learning phase compared to 82% in the decision phase. In comparison, NB and Take Two, which always predicted the same outcomes, predicted 88% of the inferences in

Table 3
Mean Proportions of Inference Outcomes and Search Processes Predicted by the Strategies in the Learning and Decision Phases of Experiment 1, Collapsed Across the High- and Low-Redundancy Conditions

	Learning Phase		Decision Phase	
	M	*SD*	*M*	*SD*
Outcomes				
Predicted by TTB	.79	.13	.82	.10
Predicted by NB/Take Two	.88	.09	.82	.09
Processes				
Search according to validity	.52	.37	.72	.32
Stop in accordance with TTB*	.25	.30	.61	.35

*Excluding guessing trials.

the learning phase and 82% in the decision phase. Thus, consistent with the costs hypothesis, TTB predicted more inferences when costs were introduced, while the fit of the compensatory strategies decreased. To examine the redundancy hypothesis, we compared the predicted inferences in the decision phases of the high- and the low-redundancy conditions. TTB predicted 87% of the inferences in the high-redundancy condition compared to 77% in the low-redundancy condition. NB and Take Two predicted 85% of the inferences in the high-redundancy condition compared to 79% in the low-redundancy condition. Thus, in accordance with the redundancy hypothesis, TTB was the strategy that predicted more inferences in the high-redundancy condition whereas the compensatory strategies predicted more inferences in the low-redundancy condition. However, due to the large overlap in strategies' predictions the differences in fit were small and do not allow a rigorous comparison of the strategies. Participants' information search then becomes the crucial test.

Two aspects of information search were analyzed: search order and stopping search. In the learning phase, participants could explore the environment without having to pay for information, but in the decision phase, cues were costly. In this situation, frugality should become crucial and an adaptive inference process would entail searching for cues selectively, in terms of both order and number of cues. In which order did participants search for information? In the learning phase, participants searched for cues in the order of their validities in on average 52% of all inferences compared to 72% of all inferences in the decision phase [$t(39) = 4.23$, $d = 0.67$, $p < .01$]. These results demonstrate that when costs for cues are introduced, participants responded adaptively, in accordance with the costs hypothesis, by more frequently searching for cues in the order of their validity. However, the search order does not discriminate between the different strategies. When applying TTB cues should be searched in the order of their validities. But as argued above, even when using a compensatory strategy it is best to search cues according to validity to limit information search. Therefore, the crucial test is when search stops.

We first determined the proportion of inferences in which search stopped before any discriminating cue was found. Since, in such a situation, no alternative was favored by the available information, this type of search indicated that participants made a guess about which alternative to choose. In the learning phase, search stopped in only 2% of all inferences before a discriminating cue was found compared to 11% in the decision phase. We then determined for all inferences in which participants did not guess whether the search stopped when one discriminating cue was found, as predicted by TTB, or continued, as predicted by the compensatory strategies (see Figure 2). Consistent with the costs hypothesis, the introduction of information costs induced participants to limit their search. Adherence to TTB's stopping rule increased substantially from the learning to the decision phase: In the learning phase, participants stopped their search in 25% of the nonguessing inferences when one discriminating cue was found, compared to 61% in the decision phase [$t(39) = 6.22$, $d = 0.98$, $p < .01$].

To test the redundancy hypothesis, we compared stopping behavior in the decision phases of the two redundancy conditions. As predicted, the proportion of TTB-consistent stopping was higher in the high-redundancy condition than in the low-redundancy condition (77 vs. 44%) [$t(38) = 3.37$, $d = 1.07$, $p < .01$]. Interestingly, this effect was found as early as in the first block of the decision phase, with TTB-consistent stopping in 76% of inferences in the high-redundancy condition, and in 43% in the low-redundancy condition.

Which of the two compensatory strategies predicted the information search better? Again, for all nonguessing inferences, we analyzed whether Take Two or NB predicted participants' information search better.[4] In the decision phase, search stopped in accordance with Take Two in 25% of all nonguessing inferences compared to 19% in accordance with NB [$t(39) = 3.93$, $d = 0.62$, $p < .01$]. Thus, Take Two predicted more accurately than NB when participants stopped their information search. Also, Take Two predicted information search slightly better in the low- than in the high-redundancy condition, 31% compared to 20% [$t(38) = 1.44$, $d = 0.46$, $p = .08$].

Is information search behavior a valid criterion with which to infer which strategy a participant has selected? To answer this question we examined whether the inferences participants made were consistent with their search behavior. When participants stopped searching after the

Table 4
Mean Proportions of Inference Outcomes and Search Processes Predicted by the Strategies in the Decision Phases of the High- and Low-Redundancy Conditions of Experiment 1

	High Redundancy		Low Redundancy	
	M	*SD*	*M*	*SD*
Outcomes				
Predicted by TTB	.87	.08	.77	.09
Predicted by NB/Take Two	.85	.08	.79	.09
Processes				
Stop in accordance with TTB*	.77	.32	.44	.30
Stop in accordance with Take Two*	.20	.26	.31	.21

*Excluding guessing trials.

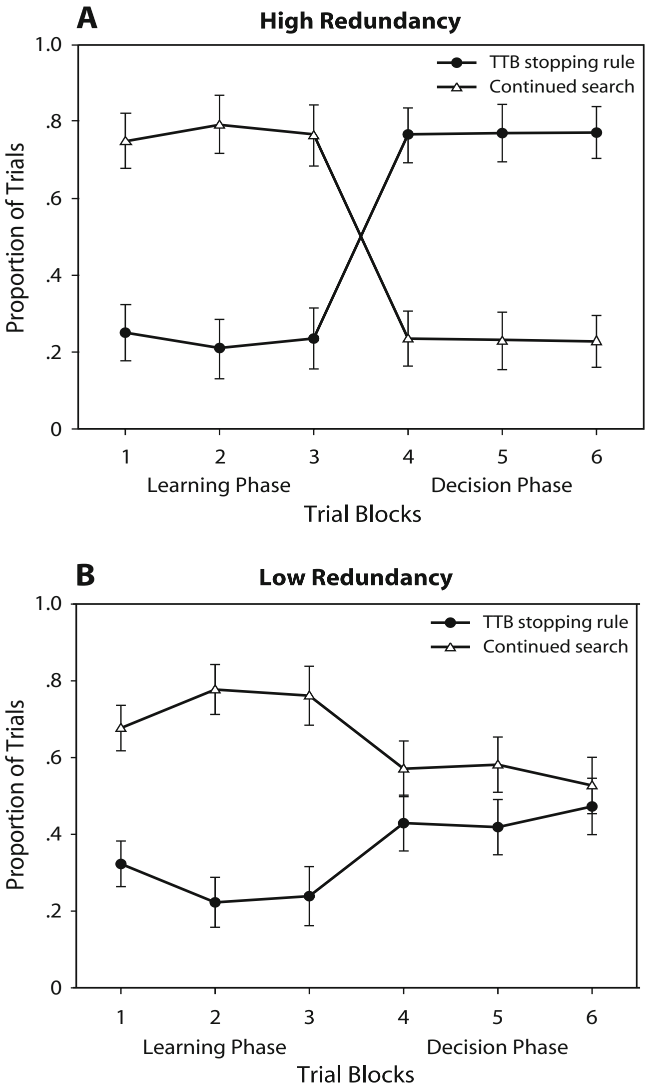

Figure 2. Proportion of nonguessing trials in which search stopped in accordance with TTB (i.e., when one discriminating cue was found) compared to the complement proportion of instances in which search continued beyond a first discriminating cue, across the six blocks of trials in Experiment 1 for (A) the high-redundancy condition and (B) the low-redundancy condition. (Error bars represent standard errors.)

first discriminating cue was found, they decided in accordance with this first cue in 99% of the cases, as predicted by TTB. In contrast, when participants searched as predicted by Take Two (or further), their inferences were in accordance with Take Two in 97% of the cases. Thus, the point at which information search stopped appears to be a valid criterion with which to infer whether people selected a noncompensatory or a compensatory strategy.

Discussion

The results of Experiment 1 provide evidence for our two hypotheses. When information became costly, participants limited their information search and more frequently searched for cues in the order of their validity. The changes in decision processes from one phase to the other occurred abruptly, while little changes could be observed within each of the two phases, supporting the interpretation that the effects are due to the introduction of costs. An alternative explanation would be that growing experience with the experimental task leads to adaptive information search, but this would be associated with a slow, continuous increase in TTB-consistent stopping.

Participants in the two redundancy conditions responded differently to the introduction of costs. In the high-redundancy condition, participants reacted mainly by stopping search as predicted by TTB. Participants in the low-redundancy environment also became more frugal when costs were introduced, but they mostly continued search beyond a first discriminating cue. Of the compensatory strategies considered, Take Two was better than NB in predicting when participants stopped their search.

How were the participants able to adapt their inference process? In the learning phase, participants in the high-redundancy condition may have experienced that there was no trade-off between accuracy and frugality: inferences based on the first discriminating cue were just as accurate as inferences that took more information into account. Participants in the low-redundancy condition could have learned that integrating information was a more accurate inference strategy than just relying on the first discriminating cue. Thus, due to outcome feedback participants could have learned to select strategies adaptively, without even noticing the degree of information redundancy. 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 mostly lead to the same decisions as using a simpler noncompensatory strategy. Since outcome feedback was given during the whole experiment, we cannot distinguish these two explanations of how participants learned to select strategies adaptively. To address this problem we conducted Experiment 2 without outcome feedback in the learning phase.

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 select when, for example, there is only limited time, or when the decision is very important. Although these situational descriptions were very general and did not include the aspect of information redundancy, the results indicate that people prefer different strategies for different situations. Along these lines, MCPL research indicated that "task feedback" alone, that is, information about the environment's structure (such as correlations between cues) without feedback about decision accuracy, can substantially improve people's inferences (Balzer et al., 1994; Balzer, Sulsky, Hammer, & Sumner, 1992).

Adaptive strategy selection based solely on task information requires that people adequately perceive the correlational structure of the environment. Research on cova-

Table 5
Mean Proportions of Inference Outcomes and Search Processes Predicted by the Strategies in the Learning and Decision Phases of Experiment 2, Collapsed Across the High- and Low-Redundancy Conditions

	Learning Phase		Decision Phase	
	M	*SD*	*M*	*SD*
Outcomes				
Predicted by TTB	.82	.13	.82	.11
Predicted by NB/Take Two	.90	.09	.82	.11
Processes				
Search according to validity	.56	.38	.72	.31
Stop in accordance with TTB*	.24	.34	.52	.34

*Excluding guessing trials.

riation assessment suggests that people might not always be good at this task, especially in the case of binary events (see Alloy & Tabachnik, 1984). However, in many of these studies the degree of covariation had to be assessed, whereas for the inference task we studied, it is sufficient to detect whether cues were correlated with each other at all, and in which direction. This seems to be a manageable task (e.g., Knowles, Hammond, Stewart, & Summers, 1972). Likewise, in Experiment 1 the detection of information redundancy might have been simplified since the cue values of all cues could be observed simultaneously for both alternatives. Thus, the frequent occurrence of divergence between cues (i.e., cues supporting different alternatives) could be used as a shortcut to identify a low-redundancy environment, whereas frequent accordance between cues (i.e., cues supporting the same alternative) is indicative of a high-redundancy environment. A precondition for detecting these characteristics is that a large amount of information is acquired, which many participants did in the learning phase of Experiment 1. Thus, it appears possible that even without outcome feedback people might be able to select inference strategies adaptively. We therefore tested the costs and redundancy hypotheses again under a condition of learning without outcome feedback.

Method

Participants. Forty participants, mostly students from various departments at the Free University of Berlin, took part in the experiment. They received performance-contingent payment of on average €10, ranging from €2 to €18. The computerized task (Czienskowski, 2004) lasted approximately 1 h.

Procedure. The task and the decision environments were identical to Experiment 1. The only aspect in which Experiment 2 differed from Experiment 1 was that no outcome feedback was provided in the learning phase. 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 meant that participants had no opportunity to learn the strategies' accuracies. Outcome feedback was only introduced in the decision phase, which was identical to the decision phase of Experiment 1.

Results

Again, the results in terms of the two hypotheses are summarized in Tables 5 and 6. First, we examined how well the strategies predicted participants' inferences. To address the costs hypothesis, we compared the two phases of the experiment. In the learning phase, NB and Take Two predicted 90% of the inferences, while TTB predicted 82%. In the decision phase, TTB and the compensatory strategies predicted the same proportion of inferences, with 82%. To address the redundancy hypothesis, we compared the decision phases of the high- and the low-redundancy conditions, finding practically no differences in the strategies' fits (in the high-redundancy condition all strategies predicted 88% of the inferences; in the low-redundancy condition, NB and Take Two predicted 76% of the inferences, compared to TTB with 75%). Due to the high overlap in the strategies' predictions, conclusions about which strategy describes the inference process best cannot be based on the percentage of predicted inferences but require the analysis of the information search process.

In which order did participants search for information? In the learning phase, participants searched according to validity in 56% of the inferences compared to 72% in the decision phase [$t(39) = 2.84$, $d = 0.45$, $p < .01$]. Thus, in accordance with the costs hypothesis, participants more frequently looked up cues in the order of their validity when costs were introduced. Since search order did not discriminate between the strategies, we focused on when search stopped.

First, the proportion of guesses was determined. In the learning phase, search stopped before a discriminating cue was found in less than 1% of inferences, while in the decision phase the proportion of guesses increased to 14%. We then determined for all nonguessing inferences whether participants stopped when one discriminating cue was found, as predicted by TTB, or whether they continued to search, as predicted by the compensatory strategies (see Figure 3). Consistent with the costs hypothesis, adherence to TTB's stopping rule increased substantially when information costs were introduced: In the learning phase, participants stopped their search in 24% of the nonguessing inferences when one discriminating cue was found compared to 52% in the decision phase [$t(39) = 5.55$, $d = 0.88$, $p < .01$].

The comparison between the decision phases of the high- and low-redundancy conditions in terms of when participants stopped their information search provided support for the redundancy hypothesis. The proportion

Table 6
Mean Proportions of Inference Outcomes and Search Processes Predicted by the Strategies in the Decision Phases of the High- and Low-Redundancy Conditions of Experiment 2

	High Redundancy		Low Redundancy	
	M	*SD*	*M*	*SD*
Outcomes				
Predicted by TTB	.88	.08	.75	.10
Predicted by NB/Take Two	.88	.09	.76	.10
Processes				
Stop in accordance with TTB*	.63	.31	.42	.34
Stop in accordance with Take Two*	.31	.21	.26	.21

*Excluding guessing trials.

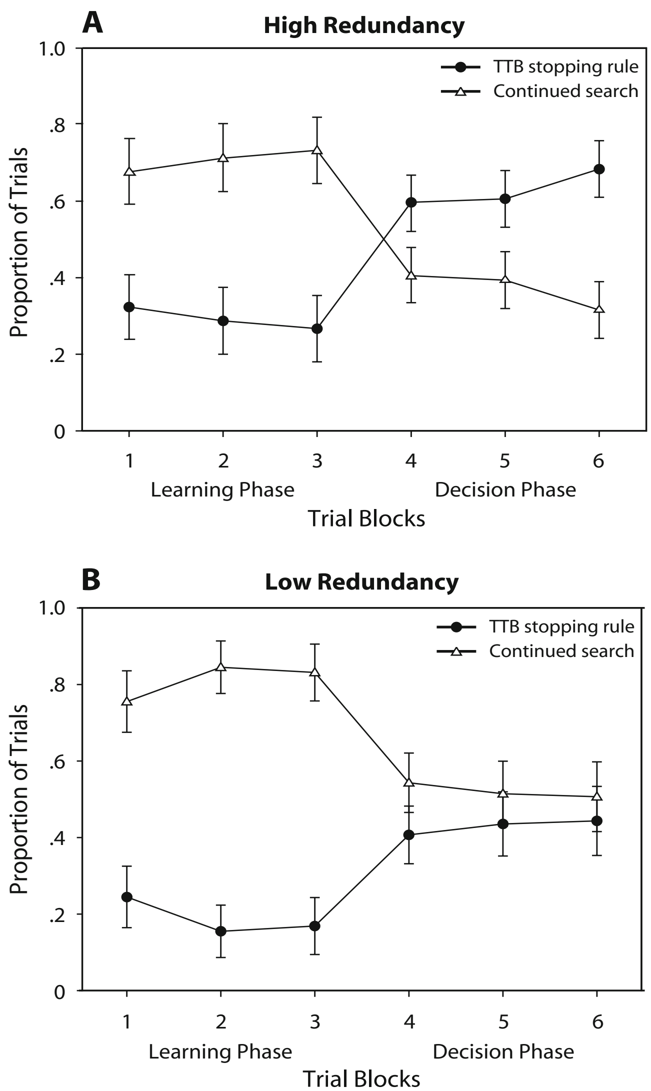

Figure 3. Proportion of nonguessing trials in which search stopped in accordance with TTB (i.e., when one discriminating cue was found) compared to the complement proportion of instances in which search continued beyond a first discriminating cue, across the six blocks of trials in Experiment 2 for (A) the high-redundancy condition and (B) the low-redundancy condition. (Error bars represent standard errors.)

of inferences in which participants obeyed TTB's stopping rule was, with 63%, on average higher in the high-redundancy than in the low-redundancy condition, with 42% [$t(38) = 1.96$, $d = 0.62$, $p = .03$]. This effect was found as early as in the first block of the decision phase, Block 4, with TTB-consistent stopping in 60% of inferences 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 [$t(19) = 2.18$, $d = 0.29$, $p = .02$].

As in Experiment 1, Take Two was better than NB in predicting when information search stopped. In the decision phase, search stopped in accordance with Take Two in 28% of all nonguessing inferences compared to 25% in accordance with NB [$t(39) = 2.30$, $d = 0.36$, $p = .01$]. In contrast to the results of Experiment 1, Take Two did *not* predict the stopping behavior better in the low-redundancy condition (with 26%) than in the high-redundancy condition (31%). But at the same time, as reported above, search continued beyond a first discriminating cue for the majority of inferences in the decision phase of the low-redundancy condition. The relatively low overall fit of Take Two (as well as NB) in the low-redundancy condition implies that participants often exhibited stopping behavior that was captured by neither Take Two nor NB.

Is the information search behavior a valid criterion with which to infer which strategy a participant has selected? Again we determined whether the search behavior was consistent with the final inference the participants made. When participants stopped searching after the first discriminating cue, they made an inference according to TTB in 99% of all cases. In contrast, when participants searched for at least as much information as predicted by Take Two, they also made an inference according to Take Two in 97% of all cases. Thus, participants' search behavior was consistent with their final inferences, illustrating that information search represents a valid criterion with which to infer selected inference strategies.

Discussion

Again, participants adapted their inference processes to both information costs and information redundancy in the environment. In the decision phase, adaptive strategy selection was observed from the first block onward. Participants mostly searched cues in the order of their validities, and in the high-redundancy condition predominantly stopped at the first discriminating cue. In the low-redundancy condition, participants mostly continued search beyond a first discriminating cue.

The observed effects in Experiment 2 were weaker than the effects observed in Experiment 1. This is presumably due to the lack of outcome feedback in the learning phase of Experiment 2. Some Experiment 2 participants might have used the newly available outcome feedback in the decision phase to evaluate the accuracy of the inference strategy that they had selected in comparison with alternative strategies, which required more extensive search. After having received confirmation that TTB was also an accurate strategy to select, participants in the high-redundancy condition could resort to a more frugal information search. In keeping with this, in the high-redundancy condition TTB's fit in predicting when participants stopped their information search increased across the last three blocks of the experiment. For the cases of continued search, Take Two was again best in predicting when information search stopped. However, many cases of continued search were captured by neither Take Two nor NB, perhaps because the class of compensatory strategies is much larger than the class of noncompensatory strategies and can include many ways in which people search for information.

In sum, Experiment 2 showed that people change their inference processes adaptively even without outcome feedback, demonstrating that the mere experience of different degrees of information redundancy in an environment can trigger the selection of adaptive strategies. The provision of outcome feedback, which enables learning of strategies' accuracies, was thus not a necessary precondition. This is an important result, because it cannot be predicted by theories that assume strategies are selected on the basis of outcome feedback, as we will discuss below. However, since the effects observed in Experiment 2 were smaller than those in Experiment 1, outcome feedback seems to enhance adaptivity.

GENERAL DISCUSSION

The first goal of our research was to test how well inference strategies perform under high versus low-information redundancy. The second goal was to test whether people select strategies adaptively in response to information redundancy. Third, inferences were examined in situations with and without outcome feedback, to test whether learning to select an adaptive strategy requires outcome feedback.

The simulation study demonstrated that the most frugal strategy in our competition, TTB, matched the accuracy of compensatory strategies in most conditions. TTB was as accurate as the other strategies in the high-redundancy environments. Astonishingly, TTB also performed as accurately as the compensatory strategies in a situation of low-information redundancy when cue validities were widely dispersed. With high dispersion of cue validities the potential advantage of information integration diminishes, because compensatory strategies will often not find enough counterevidence to overrule a previously encountered highly valid cue. Only in a situation in which both information redundancy was low and cues had similar validities did TTB's accuracy fall behind that of the competitors. As for the compensatory strategies, Take Two did as well as NB when making inferences for new, independent cases and in particular performed well in cases where TTB's accuracy was low. Thus, two simple strategies, TTB and Take Two, are sufficient to make accurate inferences under high and low-information redundancy. This is an important contribution to the rationality debate of simple heuristics (e.g., Gigerenzer & Goldstein, 1996). We show that using a simple heuristic, such as TTB, is not a second-best choice to be used only when cognitive limitations prevent the application of cognitively more demanding strategies. Instead, the adaptive selection of a simple heuristic leads to highly accurate decisions and represents rational behavior.

The experiments demonstrated that people are indeed able to select strategies adaptively in response to information costs and information redundancy. In both experiments, participants reduced their information search when information costs were introduced. More importantly, strategy selection differed systematically between the two redundancy conditions: TTB was more often selected under high-information redundancy, whereas compensatory strategies were more often selected under low-information redundancy. Finally, Experiment 2 provided evidence for adaptive strategy selection even without the opportunity to learn from outcome feedback.

Many authors have criticized that the process of how strategies are selected from a strategy repertoire is underspecified (Feeney, 2000; Luce, 2000; Morton, 2000; Newstead, 2000; Wallin & Gärdenfors, 2000). Our results add to this debate that alternative mechanisms for selecting strategies might exist, which a comprehensive theory of strategy selection will need to account for. Following Payne et al. (1993), it might be helpful to distinguish strategy selection based on the *anticipated* accuracy and costs of strategies (top-down approach) and selection based on the *experienced* accuracy and costs during the decision process (bottom-up approach). For the latter, Rieskamp and Otto (2006) have proposed a strategy selection learning (SSL) theory that specifies how people learn to select the best-performing strategies on the basis of received reinforcements, in particular outcome feedback. The results of Experiment 1 are in line with SSL, because people seemed to be able to learn, in the first phase of the experiment, to select the best-performing strategy in a given environment.

However, many inference situations do not provide outcome feedback or much learning opportunity, so people have to anticipate which strategies are likely to be successful or resort to their past experience with similar problems. According to Payne et al. (1993), strategy selection often reflects a learned contingency between certain task characteristics and a strategy's accuracy and costs. Certain environmental structures could thus serve as cues for selecting particular strategies, providing information about the likely performance of those strategies. In Experiment 2, the degree of redundancy in the environment, which could be perceived by frequent cases of accordance, or contradictions, respectively, between cues, seemed to trigger the selection of different strategies. This result cannot easily be explained by a learning theory such as SSL that assumes a learning process based on outcome feedback (for related learning theories with similar assumption see, e.g., Erev & Barron, 2005; Lemaire & Siegler, 1995). Thus, the results of Experiment 2 have important implications for such learning theories. For instance, SSL could be extended by assuming other forms of reinforcement, which could be based on past experience with similar problems or on the accordance between different strategies. Future research needs to advance these learning theories either by allowing for other forms of reinforcement or by incorporating other strategy selection mechanisms.

Our results support the notion that people select strategies adaptively in response to environmental characteristics. People's inference processes differ considerably depending on whether they encounter environments with high or low-information redundancy. Feedback about the accuracy of their decisions supports the selection of successful strategies. Yet even without outcome feedback, people adaptively select different strategies under different circumstances. This suggests that strategy selection can be driven by people's perceptions of environmental features without requiring slow, outcome-based learning in each new situation. It is an exciting enterprise to examine the origins of such adaptive strategy selection.

AUTHOR NOTE

The authors thank Richard Anderson and two anonymous reviewers for their helpful comments on earlier versions of this article, and Anita Todd and Donna Alexander for editorial assistance. Correspondence concerning this article should be addressed to A. Dieckmann, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany (e-mail: anjadieckmann@gmail.com).

REFERENCES

Alloy, L. B., & Tabachnik, N. (1984). Assessment of covariation by humans and animals: The joint influence of prior expectations and current situational information. *Psychological Review*, **91**, 112-149.

Armelius, B. A., & Armelius, K. (1974). The use of redundancy in multiple-cue judgments: Data from a suppressor-variable task. *American Journal of Psychology*, **87**, 385-392.

Balzer, W. K., Hammer, L. B., Sumner, K. E., Birchenough, T. R., Martens, S. P., & Raymark, P. H. (1994). Effects of cognitive feedback components, display format, and elaboration on performance. *Organizational Behavior & Human Decision Processes*, **58**, 369-385.

Balzer, W. K., Sulsky, L. M., Hammer, L. B., & Sumner, K. E. (1992). Task information, cognitive information, or functional validity information: Which components of cognitive feedback affect performance? *Organizational Behavior & Human Decision Processes*, **53**, 35-54.

Beach, L. R., & Mitchell, T. R. (1978). A contingency model for the selection of decision strategies. *Academy of Management Review*, **3**, 439-449.

Bergert, F. B., & Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **33**, 107-129.

Bettman, J. R., Johnson, E. J., Luce, M. F., & Payne, J. W. (1993). Correlation, conflict, and choice. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **19**, 931-951.

Brehmer, B. (1974). The effect of cue intercorrelation on interpersonal learning of probabilistic inference tasks. *Organizational Behavior & Human Performance*, **12**, 397-412.

Brehmer, B., & Joyce, C. R. B. (Eds.). (1988). *Human judgment: The SJT view*. Amsterdam: Elsevier.

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

Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 611-625.

Bröder, A., & Schiffer, S. (2003a). Bayesian strategy assessment in multi-attribute decision making. *Journal of Behavioral Decision Making*, **16**, 193-213.

Bröder, A., & Schiffer, S. (2003b). "Take the Best" versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, **132**, 277-293.

Brunswik, E. (1952). The conceptual framework of psychology. In O. Neurath (Ed.), *International encyclopedia of unified science* (Vol. 1). Chicago: University of Chicago Press.

Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, **62**, 193-217.

Christensen-Szalanski, J. J. J. (1978). Problem solving strategies: A selection mechanism, some implications and some data. *Organizational Behavior & Human Performance*, **22**, 307-323.

Chu, P. C., & Spires, E. E. (2003). Perceptions of accuracy and effort of decision strategies. *Organizational Behavior & Human Decision Processes*, **91**, 203-214.

Cooksey, R. W. (1996). The methodology of social judgement theory. *Thinking & Reasoning*, **2**, 141-173.

Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97-118). New York: Oxford University Press.

Czienskowski, U. (2004). *The oil drilling experiment* [Computer software]. Berlin: Max Planck Institute for Human Development.

Dhami, M. K., Hertwig, R., & Hoffrage, U. (2004). The role of representative design in an ecological approach to cognition. *Psychological Bulletin*, **103**, 959-988.

Dieckmann, A., & Todd, P. M. (2004). Simple ways to construct search orders. *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 309-314). Mahwah, NJ: Erlbaum.

Doherty, M. E., & Kurz, E. M. (1996). Social judgement theory. *Thinking & Reasoning*, **2**, 109-140.

Einhorn, H. J. (1970). The use of nonlinear, noncompensatory models in decision making. *Psychological Bulletin*, **73**, 221-230.

Einhorn, H. J. (1971). Use of nonlinear, noncompensatory models as a function of task and amount of information. *Organizational Behavior & Human Performance*, **6**, 1-27.

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, **112**, 912-931.

Feeney, A. (2000). Simple heuristics: From one infinite regress to another? *Behavioral & Brain Sciences*, **23**, 749-750.

Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, **103**, 650-669.

Gigerenzer, G., Todd, P. M., & the ABC Research Group (Eds.) (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.

Ginossar, Z., & Trope, Y. (1987). Problem solving in judgment under uncertainty. *Journal of Personality & Social Psychology*, **52**, 464-474.

Hammond, K. R. (1996). Upon reflection. *Thinking & Reasoning*, **2**, 239-248.

Johnson, E. J., Meyer, R. J., & Ghose, S. (1989). When choice models fail: Compensatory representations in negatively correlated environments. *Journal of Marketing Research*, **26**, 255-270.

Knowles, B. A., Hammond, K. R., Stewart, T. R., & Summers, D. A. (1971). Positive and negative redundancy in multiple cue probability tasks. *Journal of Experimental Psychology*, **90**, 157-159.

Knowles, B. A., Hammond, K. R., Stewart, T. R., & Summers, D. A. (1972). Detection of redundancy in multiple cue probability tasks. *Journal of Experimental Psychology*, **93**, 425-427.

Lee, M. D., & Cummins, T. D. R. (2004). Evidence accumulation in decision making: Unifying the "take the best" and the "rational" models. *Psychonomic Bulletin & Review*, **11**, 343-352.

Lemaire, P., & Siegler, R. S. (1995). Four aspects of strategic change: Contributions to children's learning of multiplication. *Journal of Experimental Psychology: General*, **124**, 83-97.

Luce, M. F., Bettman, J. R., & Payne, J. W. (1997). Choice processing in emotionally difficult decisions. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **23**, 384-405.

Luce, M. F., Payne, J. W., & Bettman, J. R. (1999). Emotional trade-off difficulty and choice. *Journal of Marketing Research*, **36**, 143-159.

Luce, R. D. (2000). Fast, frugal, and surprisingly accurate heuristics. *Behavioral & Brain Sciences*, **23**, 757-758.

Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory & Decision*, **52**, 29-71.

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

Mobil Oil A.G. (1997). *Erdöl und Erdgas: Suchen, Fördern, Verarbeiten* [Petroleum and natural gas: Search, extraction, refinement] [Brochure]. Hamburg: Author.

Morton, A. (2000). Heuristics all the way up? *Behavioral & Brain Sciences*, **23**, 758-759.

Naylor, J. C., & Schenck, E. A. (1968). The influence of cue redundancy upon the human inference process for tasks of varying degrees of predictability. *Organizational Behavior & Human Performance*, **3**, 47-61.

NEWELL, B. R., & SHANKS, D. R. (2003). Take the best or look at the rest? Factors influencing "one-reason" decision making. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 53-65.

NEWELL, B. R., WESTON, N. J., & SHANKS, D. R. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone "takes-the-best." *Organizational Behavior & Human Decision Processes*, **91**, 82-96.

NEWSTEAD, S. E. (2000). What is an ecologically rational heuristic? *Behavioral & Brain Sciences*, **23**, 759-760.

PAYNE, J. W. (1976). Task complexity and contingent processing in decision processing: An information search and protocol analysis. *Organizational Behavior & Human Decision Processes*, **16**, 366-387.

PAYNE, J. W., BETTMAN, J. R., & JOHNSON, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 534-552.

PAYNE, J. W., BETTMAN, J. R., & JOHNSON, E. J. (1993). *The adaptive decision maker*. Cambridge, England: Cambridge University Press.

PITT, M. A., & MYUNG, I. J. (2002). When a good fit can be bad. *Trends in Cognitive Sciences*, **6**, 421-425.

RAPOPORT, A., & WALLSTEN, T. S. (1972). Individual decision behavior. *Annual Review of Psychology*, **23**, 131-176.

RIESKAMP, J. (2006). Perspectives of probabilistic inferences: Reinforcement learning and an adaptive network compared. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **32**, 1355-1370.

RIESKAMP, J., & HOFFRAGE, U. (1999). When do people use simple heuristics, and how can we tell? In G. Gigerenzer, P. M. Todd, & the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 141-167). New York: Oxford University Press.

RIESKAMP, J., & HOFFRAGE, U. (in press). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica*.

RIESKAMP, J., & OTTO, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, **135**, 207-236.

SCHMITT, N., & DUDYCHA, A. L. (1975). Positive and negative cue redundancy in multiple cue probability learning. *Memory & Cognition*, **3**, 78-84.

SMEDSLUND, J. (1955). *Multiple-probability learning: An inquiry into the origins of perception*. Oslo: Akademisk Forlag.

SVENSON, O. (1979). Process descriptions of decision making. *Organizational Behavior & Human Performance*, **23**, 86-112.

WALLIN, A., & GÄRDENFORS, P. (2000). Smart people who make simple heuristics work. *Behavioral & Brain Sciences*, **23**, 765.

NOTES

1. Take Two has the interesting property that it can predict intransitive choices. For instance, imagine a set of three alternatives, A, B, and C. Each alternative is described by six cues, 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 alternatives A and B, Take Two selects A, because it is favored by the second and sixth cues. When comparing alternatives B and C, Take Two selects B, because it is favored by the first and third cues. However, when comparing alternatives A and C, Take Two selects C, because it is favored by the fourth and fifth cues, leading to an intransitive circle.

2. Participants were told that for 100 pair comparisons in which one alternative had a positive and the other alternative a negative cue value, the success probability (i.e., validity) specified the percentage of cases in which the alternative with the positive cue value also had the larger criterion value.

3. For instance, switching the values for cue 1 between two pair comparisons, represented by their cue difference vectors [1 1 0 1 0 0] and [−1 0 −1 −1 0 1], results in the new pair comparisons [−1 1 0 1 0 0] and [1 0 −1 −1 0 1], which decreases the average intercue correlation.

4. The overlap of Take Two's and NB's predictions of when the search process should stop is very high; they predict that search should stop at the same cue in 82% of all inferences.

APPENDIX

A hill-climbing search method was used to create environments that possessed the structure required for the simulations. It started, for each of the four conditions, with 500 artificial environments, each consisting of 50 objects, and six cues with 25 positive and 25 nonpositive cue values. The cues were ordered according to their validities. The hill-climbing search method was applied in two stages:

Stage 1 was conducted to search for environments that fulfilled the preset cue validity criteria. For all 50 objects the cue value of each object was switched with the cue value of another randomly selected object for the same cue, resulting in 50 new environments that differed from the original environment only by the two switched cue values. This was done for all cues, resulting in 300 new environments in total. Out of these 300 environments, the one with cue validities closest to the required cue validities was selected. Starting with this selected environment, again 300 new environments were created. This step was repeated until the required cue validities were obtained.

Stage 2 was conducted to search for environments with high and low-information redundancy. The search process was analogous to the process in Stage 1, the only difference being that environments were selected that either increased or decreased the average correlation between cues. If the new environments distorted the cue validities, the cue values of all six cues for two randomly selected objects were interchanged (thereby keeping the correlation fixed). This was done repeatedly until the required validities were obtained again. This step was reiterated until the average correlation between cues converged at a particular high or low average correlation between the cues.

(Manuscript received August 15, 2005;
revision accepted for publication January 10, 2007.)