

The importance of learning when making inferences

Jörg Rieskamp *

Max Planck Institute for Human Development

Abstract

The assumption that people possess a repertoire of strategies to solve the inference problems they face has been made repeatedly. The experimental findings of two previous studies on strategy selection are reexamined from a learning perspective, which argues that people learn to select strategies for making probabilistic inferences. This learning process is modeled with the strategy selection learning (SSL) theory, which assumes that people develop subjective expectancies for the strategies they have. They select strategies proportional to their expectancies, which are updated on the basis of experience. For the study by Newell, Weston, and Shanks (2003) it can be shown that people did not anticipate the success of a strategy from the beginning of the experiment. Instead, the behavior observed at the end of the experiment was the result of a learning process that can be described by the SSL theory. For the second study, by Bröder and Schiffer (2006), the SSL theory is able to provide an explanation for why participants only slowly adapted to new environments in a dynamic inference situation. The reanalysis of the previous studies illustrates the importance of learning for probabilistic inferences.

Keywords: inferences, strategy selection, heuristics, learning theory, reinforcement learning, cognitive modeling.

1 Introduction

How do people make probabilistic inferences, such as inferring the selling price of a car, the severity of an illness, or the likely winner of a tennis match? Different strategies can be applied to make these inferences, such as integrating all available information to predict the criterion. In fact, many researchers have argued that people are equipped with a repertoire of different cognitive strategies for making judgments and decisions (e.g., Brown, 1995; Brown, Cui, & Gordon, 2002; Einhorn, 1970; Fishburn, 1980; Gigerenzer, Todd, & the ABC Research Group, 1999; Ginossar & Trope, 1987; Payne, 1976; Payne, Bettman, & Johnson, 1988, 1993; Rapoport & Wallsten, 1972; Rieskamp & Hoffrage, 1999, 2008; Svenson, 1979).

Do people apply different strategies for solving probabilistic inferences? And if so, how do they select from their strategy repertoire? The cost-benefit approach to strategy selection argues that people trade off the strategies' anticipated costs and benefits. In contrast, I will argue that selection is achieved via learning—that is, people learn the success and failure of strategies through ex-

perience and select a strategy based on past success. I will describe a computational theory that specifies this learning process (Rieskamp & Otto, 2006) and will use it to explain people's probabilistic inferences in two previous studies. The goal of this article is to demonstrate that probabilistic inferences are strongly influenced by learning and that this learning process can be conceptualized as learning to select strategies.

1.1 Cost-benefit approach to strategy selection

The contingency model of Beach and Mitchell (1978) conceptualizes the selection of strategies as a cost-benefit analysis. Each strategy is assumed to lead with a specific probability to a correct solution with a beneficial outcome, and with the remaining probability to an incorrect solution with a less beneficial or detrimental outcome. Each strategy also involves some costs and it is assumed that the probability of a correct solution is positively correlated with those costs. Subtracting a strategy's costs from its benefits results in the net benefit, and the strategy with the maximum net benefit is selected by the decision maker.

Many authors have argued for such a cost-benefit approach to describe how strategies are selected (see Christensen-Szalanski, 1978; Payne et al., 1988, 1993; Smith & Walker, 1993). Yet despite being conceptually straightforward, the approach has not been spelled out as a computational model that defines how the trade-off pro-

*I would like to thank Ben Newell and Arndt Bröder for providing me with the opportunity to analyze the data of their experiments and for their helpful comments concerning the design and analysis of their studies. I also thank Anita Todd for editing a draft of this article. Correspondence concerning this article should be addressed to Jörg Rieskamp. Address: Jörg Rieskamp, Max Planck Institute for Human Development, Lentzeallee 94, 14195 Berlin, Germany. Email: rieskamp@mpib-berlin.mpg.de.

cess is cognitively determined. The first obvious barrier to a computational description consists in the fact that several benefits (e.g., monetary gains, accuracy) and costs (e.g., cognitive effort, time) can be distinguished, yet they have to be mapped onto one single scale. Second, it is not clear which strategy is applied to evaluate the costs and benefits, nor how this strategy is selected, potentially leading to an infinite regress problem of strategy selection.

1.2 Learning approach to strategy selection

Rieskamp and Otto (2006, see also Rieskamp, 2006a) took an alternative, bottom-up approach to strategy selection that is based on learning. They proposed the strategy selection learning (SSL) theory, which assumes that people most likely select the strategy that they expect to be most successful in solving an inference problem. However, instead of assuming that people deliberately trade off strategies' costs and benefits, the theory states that the strategies' *expectancies* are the result of a learning process. When facing a decision situation for the first time, people have initial expectancies for each strategy they might use, based on past experience with similar decision situations. These expectancies are updated depending on whether the selected strategy succeeds or fails to solve the decision problem. On the basis of this learning process the theory predicts that after sufficient experience an individual is most likely to select the strategy that performs best in a specific environment. Consistent with this basic assumption of the SSL theory, previous research has provided substantial empirical evidence that depending on the statistical properties of environments, different cognitive models are best at predicting people's behavior, in particular when substantial outcome feedback is provided (e.g., Bröder, 2003; Garcia-Retamero, & Rieskamp, in press; von Helversen & Rieskamp, 2008; Rieskamp, 2006a). It should be emphasized, however, that the cognitive strategies people are "selecting" are unobservable and can only be inferred from the observed behavior, such as information search or the final choices. Thus, when a strategy is capable of predicting a person's information search and final choices, this can be interpreted as indicating the person had selected the strategy. However, this is always just an interpretation and alternatively cognitive models that do not assume any cognitive strategies could provide a better account of the inference process. Hence, when herein participants' cognitive processes are described, for simplicity, as selecting specific strategies, this is always just an interpretation of a strategy's good fit in describing the observed behavior.

Although the SSL theory follows a bottom-up approach to strategy selection, it is not necessarily to be seen in opposition to the cost-benefit approach. The

trade-off process assumed by the cost-benefit approach could influence the initial preferences for specific strategies assumed by the SSL theory. However, the SSL theory argues that individuals' initial preferences for specific strategies are wiped out by the strategies' success or failure when repeatedly applied, so that over time the selected strategy for a decision problem is a function of the strategies' success. Proponents of the cost-benefit approach have also suggested that the selection process could be triggered by learning (Payne et al., 1993).

The learning approach predicts that the probability with which a specific strategy is selected for a decision problem will change over time when outcome feedback is provided. In particular, when an individual initially selects an unsuccessful strategy, the individual should switch to another more successful strategy during the course of learning. I will call this the *learning prediction*, and it implies that in a repeated decision-making situation, the strategy that most successfully predicts the majority of a person's decisions might not be the only strategy this person has applied during the course of learning. Instead, it may be that only after some time did a preference for a specific strategy develop. Therefore according to the learning approach, which strategy will predict the majority of inferences strongly depends on the speed of learning and the provided learning opportunities. When people change the strategies they select on the basis of learning, then these learning processes need to be taken into account in any accurate description of the cognitive processes underlying inferences.

1.3 The SSL theory

In the following, the SSL theory is described in more detail. Individuals have a set S of N cognitive strategies. An individual's preference for a particular cognitive strategy i is expressed by positive expectancies $q_t(i)$, so that the probability of selecting strategy i at trial t is defined by

$$p_t(i) = \frac{q_t(i)}{\sum_{j=1}^N q_t(j)} \quad (1)$$

with j as an index for the cognitive strategies. The strategies' expectancies in the first period of the task can differ and are defined by

$$q_1(i) = r_{\text{correct}} \cdot w \cdot \beta_i \quad (2)$$

where r_{correct} is the payoff received for a correct decision, w is the *initial association* parameter, and β is the *initial preference* parameter. The payoff r_{correct} received for a correct decision in a particular task is a scaling constant that allows comparisons across tasks with different payoffs. The initial association parameter w is restricted to $w > 0$ and expresses an individual's initial association

with the available strategies relative to later reinforcement, and thus it essentially describes the learning rate. The theory assumes that individuals have initial preferences for selecting particular strategies at the beginning of a task. The initial preference parameter β_i for each strategy i is restricted to $0 < \beta < 1$ and $\sum_{i=1}^N \beta_i = 1$. After a decision is made, the strategies' expectancies are updated by

$$q_t(i) = q_{t-1}(i) + I_{t-1}(i)r_{t-1}(i) \tag{3}$$

where $I_{t-1}(i)$ is an indicator function and $r_{t-1}(i)$ is the reinforcement. The reinforcement of a strategy is defined as the payoff $r_{t-1}(i)$ that the strategy produced. The indicator function $I_{t-1}(i)$ equals 1 if strategy i was selected and equals 0 if the strategy was not selected. For the following two studies it is assumed that a strategy was selected if the choice coincides with the strategy's prediction. When two or more strategies make the same prediction that coincides with the individual's choice, it is assumed that $I_{t-1}(i)$ equals the probability with which the model predicts the selection of these strategies. By definition, if $q_t(i)$ falls below a minimum value ρ due to negative payoffs, $q_t(i)$ is set to ρ ; for the following studies $\rho = 0.0001$ was used.

Finally, the SSL theory assumes that people make errors when applying a strategy, so that, by mistake, they deviate from the strategy's prediction. Let $p(a|i)$ denote the probability of choosing alternative a out of the set of alternatives when strategy i is selected, so that the probability of choosing alternative a given strategy i and an application error ε is

$$p_t(a|i, \varepsilon) = (1 - \varepsilon) \cdot p_t(a|i) + \frac{\varepsilon}{k-1} \cdot p_t(\bar{a}|i) \tag{4}$$

where $p_t(\bar{a}|i, \varepsilon)$ denotes the probability of choosing any other alternative than a from the available k alternatives, given strategy i was selected. For simplicity, the application error is assumed to be the same across strategies (for psychologically more plausible error concepts see Mata, Schooler, & Rieskamp, 2007; Rieskamp, Bussemeyer, & Mellers, 2006). The probability of choosing alternative a depends on the probabilities of selecting the strategies and the corresponding choice probabilities of the strategies, so that

$$p_t(a) = \sum_{i=1}^N p_t(i) \cdot p_t(a|i, \varepsilon) \tag{5}$$

The SSL theory is similar to other recent learning models that assume a learning process of strategy selection (see Bussemeyer & Myung, 1992; Erev & Barron, 2005; Siegler & Shipley, 1995; Stahl, 1996). These "cousins" differ from the SSL theory in the exact learning mechanisms they assume.

2 Study 1: Learning in Stable Environments

The first study that I have reanalyzed was conducted by Newell, Weston, and Shanks (2003). The authors studied the problem of inferring which of two objects had a higher criterion value on the basis of several cues. For this inference problem Gigerenzer and Goldstein (1996) proposed a simple lexicographic heuristic called Take The Best (TTB), which only considers the most valid cue for making an inference; if this cue does not discriminate the second most valid cue is considered, and so on. Gigerenzer and Goldstein illustrated through computer simulation that the heuristic performs as well as and sometimes even better than more complex alternative strategies, among them a weighted additive strategy (WADD). The WADD strategy computes a score for each alternative by taking the sum of the cue values multiplied by the cues' validities and finally selects the alternative with the largest sum.

Newell et al. argued that despite the psychological plausibility of simple heuristics it is necessary to demonstrate "that people do indeed use these heuristics in the environments in which they are claimed to operate" (2003, p. 83). For the empirical test of the heuristic they conducted two experiments, of which I have reconsidered the first one. In this experiment, the information search preceding participants' inferences led to high monetary costs. Due to these high costs it should have been obvious from the beginning of the experiment that a strategy that requires a lot of information would perform badly. Therefore when following the cost-benefit approach to strategy selection, participants should have selected an information-frugal strategy right from the beginning of the experiment. In contrast, a learning approach predicts that a well-performing strategy that requires little information will be selected more frequently after gaining experience with the inference problem.

2.1 Procedure

Newell et al.'s (2003) participants repeatedly had to infer which stock shares of a company would perform most profitably in the future. Each share was described by six dichotomous cues, with validities of the cues ranging between .65 and .90. With six dichotomous cue values for six cues, 64 different cue profiles result, leading to 2,016 possible pair comparisons, of which 180 were selected randomly for each participant. The information of the cues could be acquired with a cost of 1 pence (U.K.) for each cue. After the participants made an inference they received feedback on whether their decision was correct and earned 7 pence for each correct decision. The participants were not informed about the cues' different va-

validities. However, to facilitate learning the objective cue validities, after 60 and 120 trials the rank order of the cues according to the validities was provided to the participants.

The imposed information search cost of 1 pence for each cue was very high relative to the gain of 7 pence for a correct inference. A random choice strategy with no search costs would have led to an expected gain of 3.5 pence. Thus, by acquiring all available information it was in principle not possible to outperform the random choice strategy, even when the acquired information enabled 100% accuracy leading to a payoff of only 1 pence. This clearly illustrates that if participants anticipated the strategies' costs and benefits from the beginning of the experiment as predicted by the cost-benefit approach, they should have selected an information-frugal strategy right from the beginning.

2.2 Results

Newell et al. (2003) stated that participants' inferences were influenced by a learning process, so that they only examined the last third of the experiment to infer the strategies participants selected. To classify participants' inference strategies, Newell et al. looked at participants' search behavior and their final choices. The authors reported that at the end of the experiment the vast majority of participants searched for the cues in the order of their validities, in particular acquiring the most valid cue first. This search behavior is consistent with TTB, but even for a strategy such as WADD, search according to the cues' validities is plausible. Due to the high search costs, it is reasonable to assume that participants with a preference for WADD would try to restrict their search. This can be achieved by comparing the alternatives cue-wise, determining the difference of the cue values weighted by the cues' validities, and updating the difference with each new cue considered. Search stops whenever the present difference cannot be reversed by the outstanding cues not considered yet. With this implementation of WADD, cues would also be searched according to their validity, so that search behavior does not allow unambiguous identification of selected strategies.

For this reason it is more revealing to look at when participants stopped their information search (see also Dieckmann & Rieskamp, 2007). One third of the participants stopped search in a manner consistent with TTB, that is, they always stopped search after they found a discriminating cue. All other participants stopped their information search in 70% of all inferences after they found the first discriminating cue. In addition to looking at participants' information search, Newell et al. (2003) also considered whether TTB could predict people's choices. For 50% of the participants, TTB predicted their choices

perfectly, whereas for the remaining 50% TTB could still predict on average 80% of the choices. In sum, Newell et al. concluded that 11 of the 24 participants most likely selected TTB, whereas 9 participants used a "weight of evidence" strategy, such as WADD.

2.3 Reanalysis: Estimating the SSL theory's parameters

Newell et al. (2003) reported that the participants improved their inferences over time. Because they found substantial learning effects, they focused their analysis on the behavior at the end of the experiment. Could the strategies the participants selected at the end of the experiment be the result of a learning process, and could this learning process be described by the SSL theory? To answer this question I reanalyzed participants' behavior for the whole experiment. The parameters of the SSL theory were estimated separately for each individual's learning data as follows: The model predicted the probability with which a participant would choose each of the available alternatives for each trial conditioned on past choices and feedback. As a goodness-of-fit criterion, the G^2 measurement was used (Burnham & Anderson, 1998), defined in Equation 6, for which the likelihood function $f(y|\theta, t-1)$ denotes the probability of choice y in trial t given the model's parameter set θ and all information from the preceding trial $t-1$:

$$G^2 = -2 \cdot \sum_{t=1}^t \ln(f(y|\theta, t-1)) \quad (6)$$

The parameter values that minimized G^2 were searched for; the initial association parameter was restricted to $1 \leq w \leq 100$.

Three parameters were estimated, resulting in an average estimated association parameter value of 18. The association parameter determines the strategies' initial expectancies, and the larger the value the more time it will take to develop a preference for another strategy. A value of 18 implies that the payoff for a correct decision of 7 pence is multiplied by 18, and this "stock of association" is then divided between the two strategies the decision maker considers (cf. Equation 2). Thus, if the decision maker showed a strong preference for one strategy by, for instance, dividing up the stock with a ratio of 15 to 3, then this decision maker would need to apply the less-preferred strategy successfully at least 12 times to develop an equal preference for both strategies.

In fact, the initial preferences were not that different. The average estimated initial preference parameter was $\beta_{\text{TTB}} = .41$ for TTB, implying an initial preference of $\beta_{\text{WADD}} = 1 - \beta_{\text{TTB}} = .59$ for the WADD strategy. Thus, on average, the decision makers had a slight preference

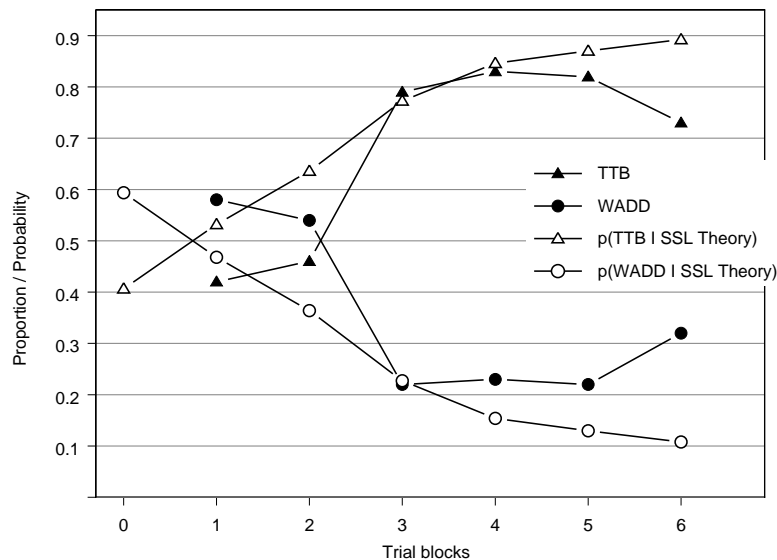


Figure 1: The proportion of choices predicted by Take The Best (TTB) and a weighted additive strategy (WADD) in the first experimental study by Newell et al. (2003). Also shown is the strategy selection learning (SSL) theory's predicted probability of selecting the two strategies. Only discriminating trials where the strategies make different predictions are shown, and each trial block consists of 30 trials.

to integrate the available information at the beginning of the experiment. This result is surprising if a cost-benefit approach to strategy selection is being followed: If participants were considering the high search costs right from the beginning of the experiment they should have preferred a strategy of focusing on single cues from the beginning. However, there were large individual differences: Although for the majority of participants ($n = 14$) an initial preference for WADD was found (i.e., $\beta_{WADD} > .50$), for a substantial number of participants ($n = 10$) an initial preference for TTB was indeed estimated (i.e., $\beta_{TTB} > .50$).

The third parameter estimated was the application error, for which an average value of .18 resulted, which is relatively high in comparison to average application errors ranging between .05 and .07 estimated across several experiments by Rieskamp and Otto (2006). A high application error implies, according to the SSL theory, that the participants often deviated from their strategies. For the experiment of Newell et al. (2003) these deviations can be easily explained by the cue validities that are required by the strategies. To determine the SSL theory's prediction, I used the objective cue validities. However, this is, of course, only an approximation, because the participants had to learn the validities during the experiment and only hints about the rank order of the validities were provided. Thus, when the subjective cue validities differ from the objective cue validities this can easily explain why a participant selecting TTB would frequently deviate from the prediction of TTB using the objective cue

validities. Yet for only four participants was a relatively high application error larger than .30 estimated. These participants apparently made their inferences more or less randomly. Newell et al. also concluded from participants' search behavior that four participants made their inferences randomly, and consistently for the SSL theory high application errors were estimated for three of these same four participants.

How well did the SSL theory describe participants' inferences? The SSL theory was able to predict participants' choices with an average probability of .71 (with $SD = 0.05$), which is slightly lower than, for instance, that found in the similar studies by Rieskamp and Otto (2006), where the choices were predicted with average probabilities of .74, .75, and .79. More interesting is whether the SSL theory can also describe the adaptive selection of cognitive strategies. The percentage of choices predicted by TTB and WADD, respectively, can be taken as an approximation of participants' strategy selections and can be compared to the probability with which the SSL theory predicts this selection per trial block.

Figure 1 shows the percentage of inferences that are predicted by TTB and by WADD, restricted to those trials where the strategies make different predictions. In addition, the figure shows the SSL theory's predicted probabilities of selecting TTB and WADD. Newell et al. (2003) provided feedback starting with the very first inference, so that participants' initial strategy preferences could not be shown. The fits of the strategies in the first trial block are already affected by learning. Therefore, the figure

additionally shows the initial probability with which the SSL theory predicts the selection of the strategy at the beginning of the experiment before any inference has been made (i.e., trial block 0).

Overall, the SSL theory's predicted probabilities of selecting the two strategies nicely matched the proportion of predicted inferences by the two strategies. Only at the beginning of the experiment does the SSL theory describe a faster learning process than is actually observed; that is, the probability of selecting TTB is larger than one would infer from the percentage of inferences predicted by TTB. This is presumably because the SSL theory makes use of the objective cue validities that had to be learned. The process of learning the validities apparently slowed down the process of learning to select an adaptive strategy. However, after the second trial block participants were familiar with the objective rank order of the cues' validities and the strongest learning process was observed, for which the SSL theory provides a good account.

2.4 Discussion

The reanalysis of the data of Newell et al. (2003) illustrates two points: First, the cost-benefit approach to strategy selection does not easily explain the experimental findings when assuming that people anticipate strategies' costs and benefits when encountering an inference problem. Second, the selection of strategies is strongly influenced by a learning process.

Due to high search costs, the cost-benefit approach predicts that participants will favor strategies that rely on little information. Contrary to this prediction, at the beginning of the experiment the WADD strategy, which requires a lot of information, predicted participants' inferences best. Thus, apparently people did not select a well-performing strategy from the beginning. Instead, the selection of strategies appears to have been strongly influenced by a learning process, so that only at the end of the experiment was the most appropriate strategy for the problem (i.e., TTB), on average, the best strategy for predicting participants' inferences.

The SSL theory was suitable to describe the observed learning process. However, because no subjective validities were elicited in the study by Newell et al. (2003), the predictions of the SSL theory were based on the objective cue validities. To the extent that these objective validities do not correspond to the subjective ones, the SSL theory will have trouble predicting the inferences. Overall, when analyzing the data of the whole experiment, it becomes clear that the behavior observed at the end of the experiment — on which Newell et al. based their conclusions — was strongly influenced by a learning process that can be accurately described with the SSL theory.

3 Study 2: Learning in dynamic environments

The experimental results of Newell et al. (2003) illustrate that people do quite well in adapting their inference processes to the inference situation. Although some participants did search for too much information, incurring high costs for search, the majority made their inferences in accordance with simple inference processes that were most suitable for the high-search-costs situation. The reanalysis shows that the adaptive behavior Newell et al. reported for the end of the experiment is the result of a learning process that can be described by the SSL theory.

Do people also accurately adapt their inference process in a dynamic inference situation in which the best-performing strategy changes over time? This question was addressed in a study by Bröder and Schiffer (2006). In the first half of their experiment participants encountered either an environment in which it was best to focus on single pieces of information (i.e., TTB led to a better performance than WADD) or an environment in which WADD outperformed TTB. Thus, the central question was whether the participants would select the most adaptive strategy according to the environment. Moreover, for half of the participants the environment changed after the first half of the experiment, that is, they were confronted with a second, not previously encountered environment. Would the participants also be able to adapt to the new environment?

3.1 Procedure

In the first of Bröder and Schiffer's (2006) experiments the participants had to infer which of three companies' shares of stock would perform most profitably in the future. The experiment had two phases with 80 inferences each. The participants encountered either a compensatory or a noncompensatory environment in the first phase, in which WADD or TTB performed best, respectively. Thereafter, for half of the participants the environment was changed so that the environment they had not previously encountered was used for the second phase (and half of the participants for whom the environment changed got a hint that the environment had changed). The payoff that each share produced was a function of the four cues and the amount of information the participants acquired (each acquired cue value led to a reduction of the final gain of 4%). The three shares led to a specific payoff and in principle all shares could lead to a similar positive payoff. Thus, whereas in the study by Newell et al. (2003) there was one single correct choice with a gain, in the study by Bröder and Schiffer the options led to diverse outcomes and these were presented to the participants after each choice. This experimental aspect is

important, because it defines the task's incentive structure and the feedback that induces a learning process. According to the SSL theory, a strategy's reinforcement is defined by the monetary gains it produces. Therefore, if TTB and WADD lead to different choices, but the payoffs attached to these choices do not differ substantially, then this would not induce a strong learning effect in comparison to a learning situation in which only one choice is rewarded.

As in the study of Newell et al. (2003), the participants had to learn the importance of the cues and no cue validities were provided. This is important because again the strategies' fits in predicting participants' inferences were based on the objective cue validities. This implies that a low fit of TTB could mean either that the strategy was not selected at all or that the strategy was applied differently by using different subjective cue validities. To determine the fit of the compensatory strategy WADD Bröder and Schiffer (2006) assumed that the rank order of WADD's weights would correspond to the rank order of the objective validities. Furthermore, three different variants of WADD were examined, with different weighting schemes for the weights of the cues, and the variant with the highest fit was assigned to the participants. Naturally, this implies that the participants had a priori a higher chance of being classified as using WADD. To avoid this a priori advantage, in the following reanalysis I only used one single weighting schema to determine WADD's predictions. I selected the schema according to which any two less important cues were always sufficient to compensate for the information of one, more important cue. Moreover, I did not include a compensatory strategy that gave equal weight to each cue as Bröder and Schiffer did. Thus, I used only one compensatory and one non-compensatory strategy to reanalyze the data.

3.2 Results and discussion

Bröder and Schiffer's (2006) data were reanalyzed by classifying each participant as using either TTB or WADD.¹ For the first phase of the experiment a strong effect of the environment on the strategy classification was observed.

¹Classification was performed on the basis of the quadratic scoring rule (Selten, 1998), according to which a strategy i with application error ε reaches a fit QS of $QS(i, \varepsilon) = 1/T \sum_{t=1}^T \sum_{k=1}^3 [D_t(k) - p_t(k|i, \varepsilon)]^2$, with t for the trial, T as the number of trials, and k for the three alternatives; $p_t(k|i, \varepsilon)$ denotes the probability of choosing alternative k when using strategy i (see Equation 4), and $D_t(k)$ is an indicator function that equals 1 if alternative k was chosen and equals 0 if the alternative was not chosen. For each strategy and participant the optimum value for the application error ε was selected that led to the lowest, that is, the best value for QS . Participants were assigned to the strategy with the lowest QS value and were left unclassified if both strategies reached the same fit.

For the noncompensatory environment the majority of participants (79%) were classified as using the noncompensatory TTB heuristic, whereas in the compensatory environment 58% of the participants were classified as using the compensatory strategy WADD (ignoring three unclassified participants across all conditions). Thus, the participants reacted sensitively and adaptively to the two environments.

What happened in the second phase of the experiment? For those participants who continued with the same environment, results were similar to those in the first phase of the experiment; that is, 77% of the participants were classified as using TTB in the noncompensatory environment and 70% were classified as using a compensatory strategy in the compensatory environment. Thus, participants apparently just continued to make their inferences in the same way as they had in the first half of the experiment, and in the compensatory environment participants apparently selected WADD more frequently.

Most interesting is what happened when the participants encountered a new environment. Did they adaptively switch their strategy? An adaptive switch was only partly observed: Only 27% of the participants who encountered the compensatory environment in the second phase after seeing the noncompensatory environment in the first phase were classified as using a compensatory strategy. Thus, apparently the participants only slowly adapted to the new environments. However, 65% of the participants who encountered the noncompensatory environment in the second phase after seeing the compensatory environment in the first phase were indeed classified as using the better performing TTB strategy.

These results are somewhat different from the results reported by Bröder and Schiffer (2006), because in their analysis three additional compensatory strategies were included. Accordingly more participants were classified as using a compensatory strategy and fewer as using the noncompensatory strategy TTB, and in particular a switch to TTB was not observed when participants encountered the noncompensatory environment after the compensatory environment. It is an open question whether Bröder and Schiffer classified more participants as using a compensatory strategy because they simply examined more compensatory strategies, or whether the participants indeed used different types of compensatory strategies. However, this question is not of particular importance for the present article.

3.2.1 Describing the learning process with the SSL theory

Bröder and Schiffer (2006) described the lack of adaptivity in the second phase of the experiment as a routine effect; that is, they assumed that the strategies the par-

ticipants selected in the first phase of the experiment became routines, so that when the environment changed the participants stuck to their routines. To explain why participants in the second phase did not switch to another strategy they referred to a dual-system theory of decision making (e.g., Betsch, 2005; Evans & Over, 1996; Sloman, 1996). According to the dual-system approach, people can be in either of two cognitive modes when making decisions, a deliberative, top-down mode or an associative, experienced-based, bottom-up mode. Bröder and Schiffer argued that when participants encountered a novel decision situation they were in a top-down mode for “calculations of long-term payoff advantages” that led to the selection of an adaptive strategy. Thus, the top-down mode refers to the cost-benefit approach, which should lead to the selection of the best-performing strategy in an environment. In contrast, after participants selected a successful strategy they “switched to a bottom-up mode by routinely applying this strategy, which avoids testing the consequences in each trial” (p. 915). However, neither proponents of the two-system approach nor Bröder and Schiffer provided a computational model that specifies how these two modes of cognition are activated, determines how and whether they interact, and describes the observed learning processes.

In the following I will show how the SSL theory explains the results. Bröder and Schiffer (2006) also referred to the SSL theory as a potential description of the bottom-up mode of strategy selection. I will show with the reanalysis that the SSL theory is able to describe the learning process not only for the second half of the experiment but also for the whole experiment by assuming a learning process of strategy selection. Thus, instead of postulating a dual-system theory of decision making, I can illustrate that a single “cognitive mode” of learning is sufficient to explain the experimental findings. I will argue that the SSL theory provides a more parsimonious explanation of the maladaptive selection of strategies described by Bröder and Schiffer than the dual-system approach does.

The parameters of the SSL theory were estimated separately for each individual’s learning data as follows: The model predicted the probability with which a participant would choose each of the available alternatives for each trial conditioned on past choices and feedback (using maximum likelihood as a goodness-of-fit criterion, cf. Equation 6). The predictions of TTB and WADD were determined on the basis of the objective weights of the cues as described above. For the SSL theory three parameters were estimated, resulting in an average estimated association parameter value of .36; this is similar to values estimated in Rieskamp and Otto (2006). The average estimated initial preference parameter was $\beta_{TTB} = .61$ for TTB, implying an initial preference of $\beta_{WADD} = 1 - \beta_{TTB}$

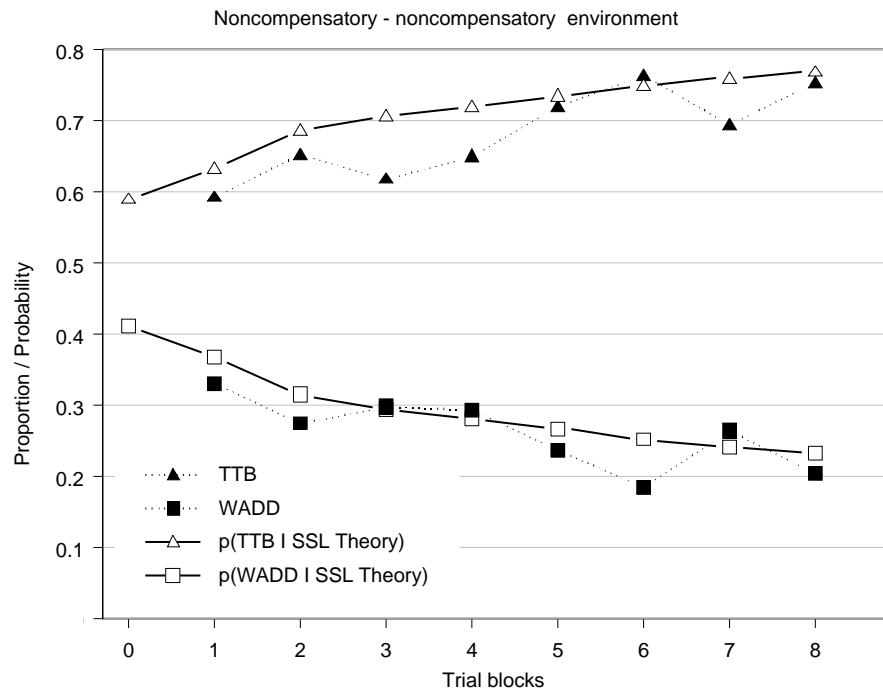
$= .39$ for WADD. Thus, on average, the participants had an initial preference for TTB. This is surprising considering that in similar studies an initial preference for WADD has often been observed. One explanation might be the high search costs. In Bröder and Schiffer’s (2006) study the participants had to pay 4% of their potential gain for every single cue value, meaning that if they searched for all the information their gain was reduced by 48%. This procedure might have made it more salient that a strategy that requires a lot of information cannot perform well, leading to an initial preference for TTB.

The third parameter estimated was the application error, for which an average value of .26 resulted. This value is relatively high in comparison to the results of Rieskamp and Otto (2006), with average application errors ranging between .05 and .07 across several experiments. A high application error according to the SSL theory implies that the participants often deviated from their selected strategies. For Bröder and Schiffer’s (2006) experiment this high value can be explained by the subjective importance the participants gave to the cues. To determine the predictions of the strategies one particular compensatory weighting schema was employed. However, this is only an approximation and the participants might have given rather different subjective importance to the cues. If the subjective importance differed from the objective weights it is not surprising that TTB or WADD relying on the objective weights do not predict all inferences, which is reflected in a high application error. Note, however, that the SSL theory does not require that people learn specific weights; it would only be preferable to elicit the subjective importance of the cues, so that the strategies’ predictions are determined on the basis of subjective importance.

How well did the SSL theory describe participants’ inferences? The SSL theory was able to predict participants’ choices with an average probability of .58 (with $SD = 0.05$), which is much larger than the .33 one would predict by random chance. Nevertheless, it is lower than the average predicted probability of .75 found in Rieskamp and Otto’s (2006) experiment, which also examined a three-alternative inference problem. The lower fit is presumably due to the employed weights for the cues that only partly correspond with the subjective importance the participants gave to the cues.

Can the SSL theory also describe the adaptive selection of cognitive strategies? Figure 2 shows for the two stable environments the percentage of choices predicted by TTB and WADD across eight trial blocks of 20 trials each, restricted to those trials where the two strategies make different predictions. Additionally, the figure shows with what probability the SSL theory predicts the selection of TTB and WADD. Figure 2A shows the results for the participants only facing a noncompensatory

A



B

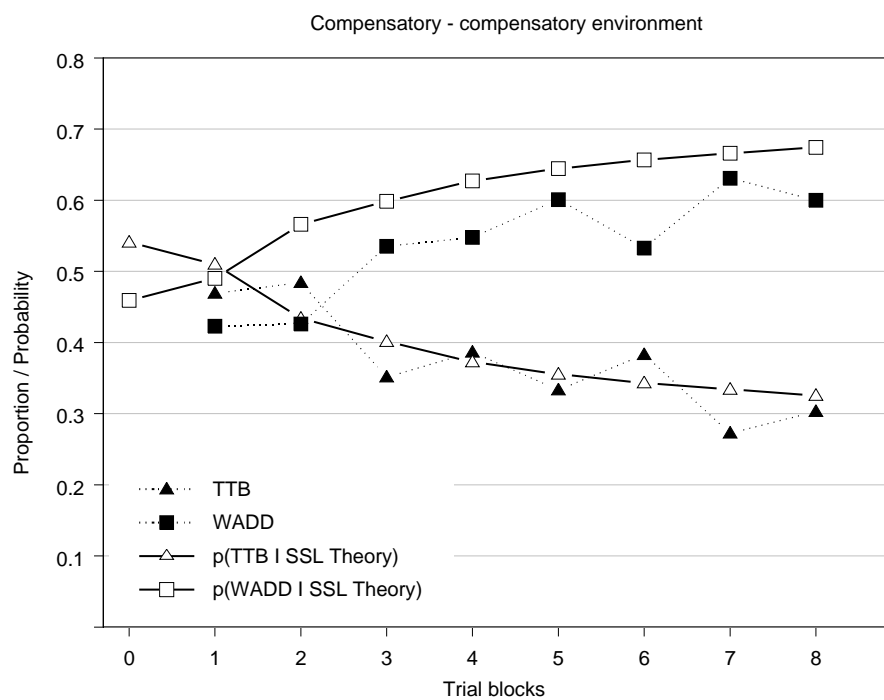


Figure 2: The proportion of choices predicted by TTb and WADD in the stable noncompensatory (A) and stable compensatory (B) environment conditions of the first experimental study of Bröder and Schiffer (2006). Also shown is the SSL theory's predicted probability of selecting the two strategies. Only discriminating trials where the strategies make different predictions are shown, and each trial block consists of 20 trials.

environment and illustrates a strong learning effect: The proportion of choices that could be best predicted by TTB steadily increases across the eight blocks of trials. This learning effect explains why the majority of participants were classified as using TTB in both halves of the experiment.

A corresponding learning effect was observed for the compensatory environment (Figure 2B). Here, in the first two trial blocks TTB predicted more inferences than WADD, but due to learning, the percentage of inferences predicted by WADD increased so that in the second half it predicted the majority of inferences. This learning effect can explain why most participants were classified as using WADD in the second half of the experiment. The compensatory environment also illustrates that the participants did not select the most adaptive strategy right from the beginning, as a cost-benefit approach predicts. Instead, a preference for WADD was developed through the first 80 choices. The probabilities with which the SSL theory predicts the selection of TTB or WADD nicely match the proportion of inferences predicted by the two strategies. These results suggest that people do not initially select a strategy for an environment and keep using it without monitoring its success, as suggested by Bröder and Schiffer (2006, see p. 915). Instead people's initial preferences for a specific strategy are quickly wiped out by the experiences they have, so that the strategy they select for a specific environment is essentially the result of the strategies' success.

3.2.2 Describing behavior in dynamic environments with the SSL theory

But do people also learn to adapt to dynamic environments? Figure 3 shows for the dynamic environments the percentage of choices predicted by TTB and WADD. Figure 3A shows the results for the experimental condition starting with the noncompensatory environment and continuing with the compensatory environment. For the first half of the experiment a learning process is observed with an increasing proportion of inferences predicted by TTB. After the shift of the environment in the fifth trial block the percentage of inferences predicted by TTB decreases only slowly. In the seventh trial block a larger decrease can be observed. Nevertheless, across all trial blocks in the second phase of the experiment when participants encountered the compensatory environment TTB did much better than WADD in predicting the inferences. Can this maladaptive behavior be described by the SSL theory?

The SSL theory predicts an even slower adaptation process to the new environment, for two reasons: First, when TTB is selected with a high probability a person will only rarely select the competing compensatory strategy and will therefore only rarely experience its better

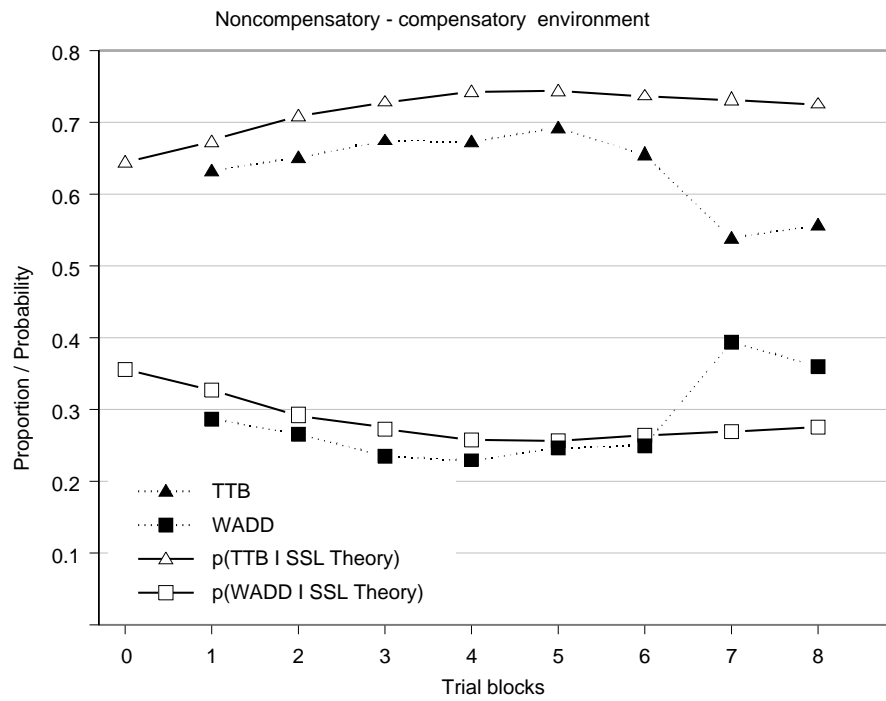
performance compared to TTB. Second, as a result of the first phase of the experiment, the strategies' expectancies, in particular TTB's expectancies, have grown substantially, so that it will take a considerable number of trials with positive reinforcement for WADD's expectancies to exceed TTB's expectancies again.

Figure 3B shows the results for the condition starting with the compensatory environment and continuing with the noncompensatory environment. Here, the results are less clear. In the first half of the experiment the percentage of inferences that could be best predicted by WADD, the better performing strategy in the compensatory environment, increased across the first four trial blocks, so that in the third and fourth trial block it predicted more inferences than TTB. In the second half of the experiment the percentage of inferences predicted by WADD decreased with a corresponding increase of the inferences predicted by TTB, so that TTB did better in the last three trial blocks. Thus, the proportion of predicted inferences by WADD and TTB changed adaptively depending on the environment. However, the differences are rather small and less conclusive, considering that the proportion of inferences predicted by TTB or WADD varies between approximately 40 and 55%. Nevertheless, the SSL theory describes this learning process by an increasing probability with which WADD is selected in the first half of the experiment, followed by a decreasing probability in the second half. Thus, the SSL theory can also explain when an adaptation to a new environment occurred, for instance, when previous experience did not lead to a strong preference for a specific strategy.

Bröder and Schiffer's (2006) experimental results illustrate that people are able to select strategies adaptively in a stable environment. However, in a dynamic environment in which the best strategy to solve an inference problem changes, people do not always switch to the better performing strategy. These results can be explained by the SSL theory. First, the selection of the best-performing strategy in the stable environment appeared to be the result of a continuous learning process, rather than of an initial cost-benefit trade-off process. In fact, the strategy that predicted the most inferences at the end of the experiment was not necessarily the strategy that predicted more inferences at the beginning, as illustrated with the compensatory environment.

The SSL theory also explains the lack of adaptivity: When people develop a strong preference for a specific strategy based on their experience, as, for instance, for TTB in the noncompensatory environment, they will too rarely select alternative strategies to be able to detect their potentially superior performance. However, this does not mean that they do not change their inferences at all. When considering the proportion of inferences across the different trial blocks it becomes clear that the participants

A



B

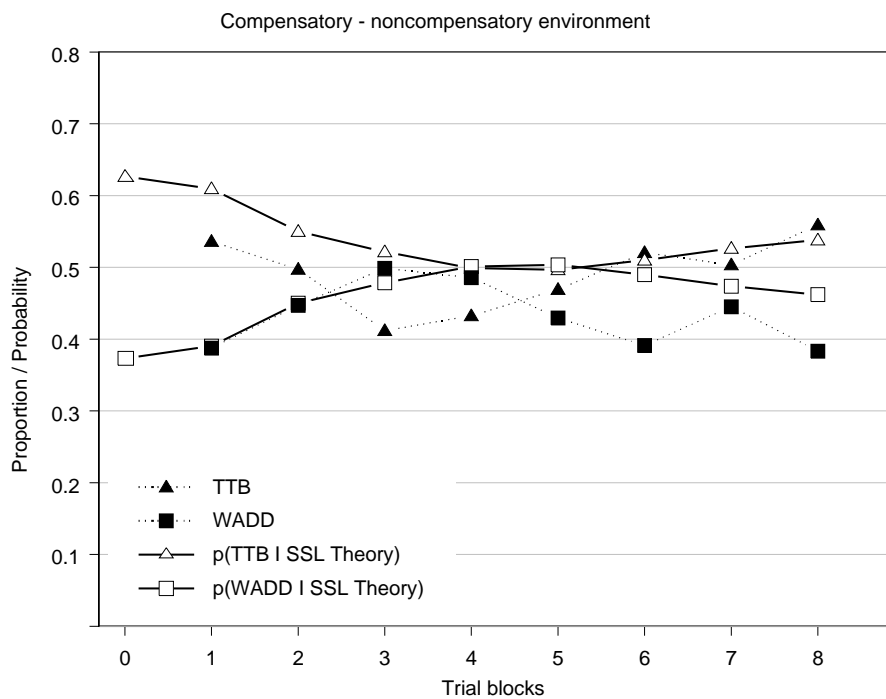


Figure 3: The proportion of choices predicted by TTb and WADD in the dynamic environment condition of the first experimental study of Bröder and Schiffer (2006), starting with the noncompensatory (A) or the compensatory (B) environment. The environment was changed after the fourth trial block. Also shown is the SSL theory's predicted probability of selecting the two strategies. Only discriminating trials where the strategies make different predictions are shown, and each trial block consists of 20 trials.

did slowly change their inferences even when they had developed a strong preference for one strategy.

3.2.3 Assuming a forgetting process during learning

The SSL theory predicts maladaptive behavior in a dynamic environment when a person has developed a strong preference for a specific strategy. However, this maladaptive behavior was less pronounced than predicted by the SSL theory (see Figure 3A). The theory predicts that an adaptation to the new environment occurs only slowly. To examine how much learning it would take before the theory predicts that the better performing strategy WADD is selected again, I simulated a hypothetical continuing learning process beyond the actual observed 160 trials. When using the estimated parameters it would take approximately 100 additional trials before the better performing strategy WADD becomes most likely to be selected (assuming an average payoff of -25 for TTB and 25 for WADD). This adaptation process appears too slow in comparison to participants' real adaptation.

How could a quicker adaptation process be accomplished? One peculiarity of the SSL theory consists in assuming no forgetting during learning, unlike in many learning models (e.g. Busemeyer, & Myung, 1992; Erev & Barron, 2005; Rieskamp, 2006b; Rieskamp, Busemeyer, & Laine, 2003). According to the SSL theory, the reinforcement a strategy received a long time ago has the same effect as a recent reinforcement, which is psychologically not very plausible. Moreover, without any forgetting, adaptations in dynamic environments occur very slowly. Thus, it appears reasonable to incorporate a forgetting process in the SSL theory; accordingly Equation 3 can be changed to:

$$q_t(i) = (1 - \varphi) \cdot q_{t-1}(i) + I_{t-1}(i)r_{t-1}(i) \quad (7)$$

with the forgetting parameter $0 \leq \varphi \leq 1$. This extended version of the SSL theory will predict a more dynamic learning process than the original model, when the forgetting parameter is below 1. For instance, a forgetting parameter value of $\varphi = .20$ implies that a reinforcement of 100 that was received 10 trials earlier will have a present value of only 10 for a strategy's expectancy, so that only the present reinforcement will strongly affect a strategy's evaluation. Thus, incorporating a forgetting process makes the SSL theory able to cope with dynamic environments. However, it also makes the theory more complex, which was the reason why the forgetting process was not incorporated in the original specification of the SSL theory (see Rieskamp & Otto, 2006).

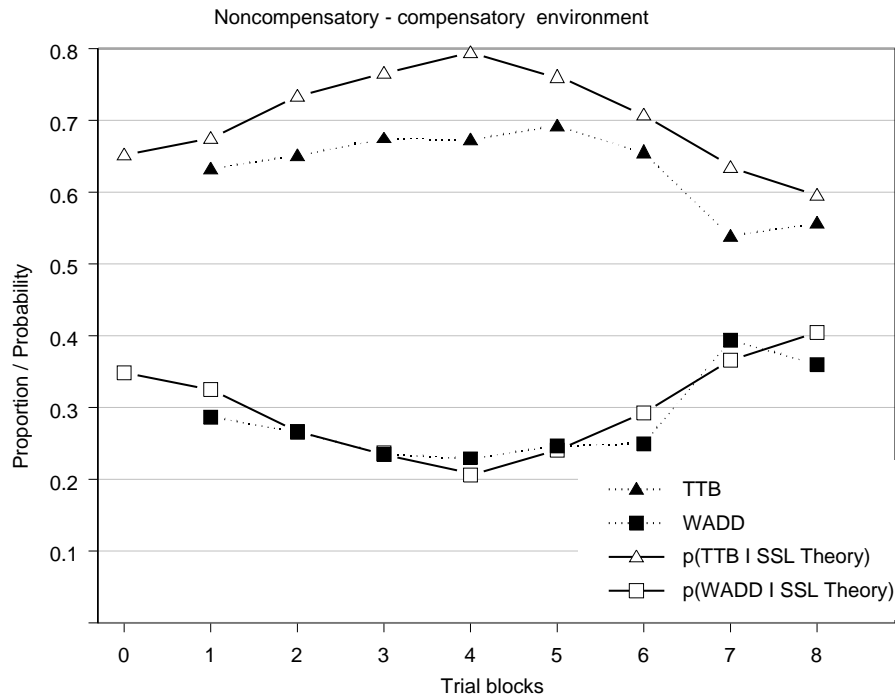
To examine whether the learning process in the study by Bröder and Schiffer (2006) could be better described by the extended SSL theory incorporating a forgetting process, I estimated the parameters of the extended SSL

theory separately for each individual's learning data as described above. For the extended SSL theory four parameters were estimated, resulting in an average estimated association parameter value of 16, an average initial preference parameter of $\beta_{\text{TTB}} = .59$ for TTB, an application error of .26, and finally an average forgetting rate of $\varphi = .07$ ($SD = .11$). Thus, the standard parameters were similar to those for the original SSL theory reported above, with the only difference being a lower value for the initial association parameter. The average forgetting rate of .07 is relatively moderate, so that previous reinforcements are not immediately forgotten but affect the selection process for some time; for instance, a reinforcement of 100 will still have a present value of approximately 50 after 10 trials.

The extended SSL theory was able to describe the slow adaptation process in the dynamic environments. Figure 4 shows the probability with which the theory predicts the selection of TTB and WADD. Especially when examining Figure 4A, representing the condition where participants faced the noncompensatory before the compensatory environment, it becomes clear that by assuming a forgetting process a more dynamic learning process will be predicted, which describes the observed behavior much better than the original SSL theory. Figure 4A shows that by assuming a forgetting process the probability of selecting TTB decreases continuously when the new environment is encountered, providing a good account of the observed slow adaptation process. Likewise when considering Figure 4B, representing the condition where participants faced the compensatory environment first, the extended SSL theory predicts a more dynamic learning process.

Although the extended SSL theory described a more dynamic learning process in line with the observed behavior it did not predict the choices with a larger probability. The extended SSL theory predicted participants' choices with an average probability of .58, which is identical to the predicted probabilities of the original SSL theory. To test which theory is more appropriate to describe the observed inferences, I determined for each participant which theory had a better Akaike information criterion (AIC; defined as $AIC = G^2 + 2 \times \text{number of free parameters of the model}$). The AIC takes a model's complexity into account by adding a penalty term to the fit of the model (for details see Burnham & Anderson, 1998). For only 30 of the 120 participants did the extended SSL theory assuming a forgetting process have a better AIC, whereas for the remaining majority of 90 participants the original SSL theory had a better AIC value (this result did not differ substantially for the four conditions examined). Thus, when taking the models' complexity into account the original SSL theory is the preferable model. Although the extended SSL theory is able to account for a more dy-

A



B

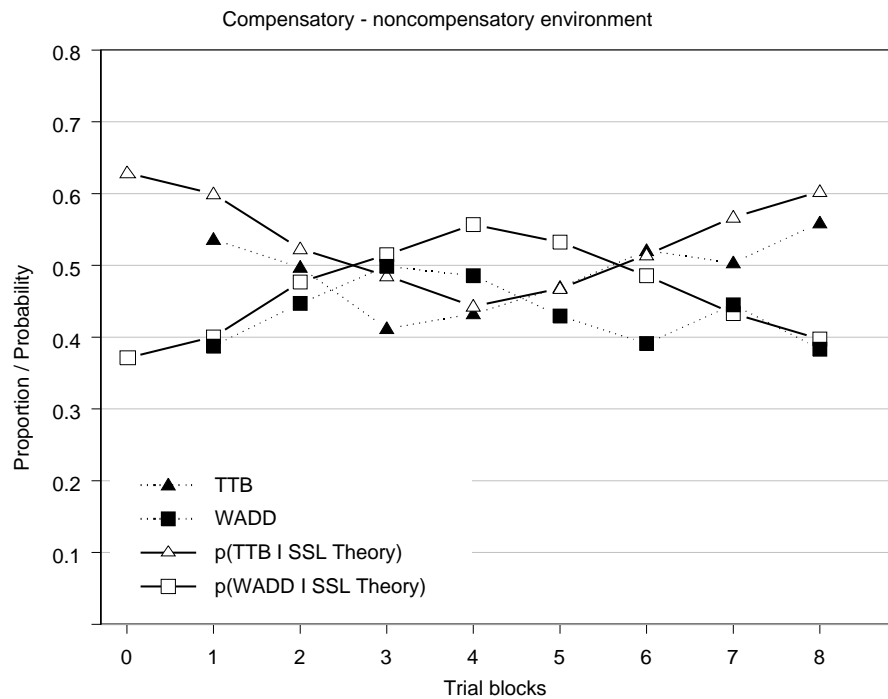


Figure 4: The proportion of choices predicted by TTb and WADD in the dynamic environment condition of the first experimental study of Bröder and Schiffer (2006), starting with the noncompensatory (A) or the compensatory (B) environment. The environment was changed after the fourth trial block. Also shown is the predicted probability of selecting the two strategies by the extended SSL theory incorporating a forgetting process. Only discriminating trials where the strategies make different predictions are shown, and each trial block consists of 20 trials.

dynamic learning process, this does not allow the theory to predict the inferences with a much higher probability.

4 General discussion

How do people select strategies for making inferences? Does learning provide a fruitful basis for explaining how strategies are selected and do people select successful strategies even in dynamic environments? These main questions of the article guide the discussion.

4.1 How strategies are selected

If people are equipped with a repertoire of strategies, how do they decide which one to follow? The most prominent answer in the literature to this question is to assume that people trade off the costs and benefits of a particular strategy and select the one with the overall best evaluation. Although this approach appears reasonable, the specifics of the trade-off process have not been computationally specified. This may be because different pathways have been proposed for how the trade-off process could take place. Payne et al. (1993) suggested that the trade-off process could be based on a deliberate evaluation of the anticipated costs and benefits of the decision strategies (pp. 92–99), yet they also argued that this process does not necessarily need to take place deliberately and could be based on learning (pp. 200–201). Both trade-off processes appear intuitively appealing, but if they are both responsible for the selection of strategies, it is necessary to specify under which circumstances each process will take place or how the two processes interact with each other.

Dual-system theories of cognition provide an alternative explanation of how people select decision strategies that is related to the cost-benefit approach: People are in either a “cognitive mode” in which the pros and cons for selecting specific strategies are deliberately traded against each other, or an “associative mode” in which they do not think deliberately about the selection process, but the selection is driven by a slow learning process. Recently dual-system theories of cognition have gained increasing attention, yet despite their popularity they face the same problem as the cost-benefit approach: Namely, they are not sufficiently specified to allow precise predictions. To make these predictions it would be necessary to specify when each cognitive system is active or dominates the other, and how these two systems operate. When the theories are computationally specified it will be possible to test the different conceptualizations of the dual-system theories rigorously against each other, and against alternative “one-system” theories to evaluate their advantages, a necessary test that has not yet been

accomplished.

I have followed a third path to explain how decision strategies are selected, one that is based on learning. The SSL theory assumes that people have initial preferences for selecting specific strategies. It does not explain where the initial expectancies come from, but they are most likely the result of previous experiences with the strategies or the results of an initial evaluation of the strategies. However, more importantly, the initial preferences are quickly changed through the success or failure of the selected strategies. Therefore the strategy a person selects after gaining some experience with the inference situation is basically determined by the performance of the strategies being considered. For both studies that were reanalyzed for the present article, the SSL theory provided an appropriate description of the selection process. Therefore I would argue that the learning assumption is a very fruitful perspective to follow when explaining how people select strategies from their repertoire.

4.2 The importance of learning when interpreting people’s inferences

When analyzing participants’ inference strategies Newell et al. (2003) restricted their main analysis to the end of the experiment. They argued that at the beginning of the experiment people would learn and explore the structure of the environment. However, due to the large search costs used in their experiment, when following the cost-benefit approach one would have expected the use of information-frugal strategies right from the beginning. Yet the reanalysis of their experiment illustrated that it took some time before the participants converged on selecting a specific strategy. When considering the results shown in Figure 1 it becomes evident that the strategies that do well in predicting participants’ inferences at the end of the experiment do not predict the majority of inferences at the beginning of the experiment. It appears that learning is an important factor that needs to be taken into account when interpreting inferences in a situation in which outcome feedback is provided. Therefore the conclusions regarding how people make their inferences depend on the provided learning opportunity. Depending on whether the learning opportunities are sufficient to allow people to adapt to the specific environment, conclusions might differ concerning whether people make their inferences adaptively.

The importance of learning for interpreting people’s inferences also becomes obvious when considering the reanalysis of Bröder and Schiffer’s (2006) first experiment. In one experimental condition with a stable environment the participants made altogether 160 inferences, without any initial learning phase and with 20 fewer inferences than participants made in Newell et al.’s (2003) experi-

ment. Bröder and Schiffer simply analyzed the data from the onset of the experiment when classifying participants' inference strategies. It is obvious that this classification was strongly influenced by the observed learning effects represented in Figure 2. For instance, when looking at the results for the compensatory environment represented in Figure 2B, it becomes clear that the compensatory strategy did best in predicting the inferences at the end of the experiment, whereas at the beginning the compensatory and noncompensatory strategies, on average, predicted the inferences equally well. Bröder and Schiffer classified their participants as using a specific strategy essentially depending on which strategy predicted the majority of choices. However, which strategy will predict the most choices depends on learning. A person who initially selects TTB in the compensatory environment might slowly adapt to the environment by selecting a compensatory strategy at the end of the experiment. However, this person might still be classified as selecting TTB, because the majority of choices are best predicted by TTB. Thus, for this person the incorrect conclusion of not selecting an adaptive strategy would be made.

The importance of learning is particularly crucial when considering behavior in a dynamic environment. In Bröder and Schiffer's (2006) second experimental condition the participants were only provided with 80 trials before the environment was changed, after which another 80 inferences were made. If people adapt only slowly to new environments, 80 trials might not be sufficient for adaptation to occur (see, for instance, the slow adaptation process illustrated by Figure 3A). However, if the participants had been given more learning opportunity, such as the 120 trials in Newell et al.'s (2003) study, this might have been sufficient to induce them to switch to selecting an adaptive strategy. These examples illustrate that as a result of experience, people change the strategy they select for making their inferences. When testing whether people select adaptive strategies it is necessary to provide sufficient learning opportunity.

4.3 The limitations of the SSL theory for strategy selection

The SSL theory provides a fruitful account of how participants select their inference strategies. For the study of Newell et al. (2003) the theory was able to account for how the predominant adaptive behavior at the end of the experiment was the result of a learning process. Likewise the SSL theory was able to describe the learning process observed in the stable environments of Bröder and Schiffer's (2006) first experiment. Moreover, the theory could explain why maladaptive behavior was observed in a dynamic environment: Maladaptive behavior occurs when a decision maker develops a strong preference for a specific

strategy. As a consequence, alternative strategies are too rarely selected to "discover" their potential performance advantage after an environment change. Furthermore, according to the SSL theory, reinforcement for a strategy is accumulated over time, so it will take a while before a better performing strategy can accumulate enough reinforcements to make its selection most likely.

Yet it became clear that the learning process predicted by the SSL theory was too slow compared with the observed learning process. To accommodate a faster learning process the SSL theory can be extended to include a forgetting process, so that strategies' expectancies decline over time if they do not receive reinforcement. With such a forgetting process a more dynamic learning process occurs, implying a quicker adaptation to a new environment, as illustrated in Figure 4. However, the increased complexity of the extended SSL theory incorporating forgetting could not be justified with a substantially higher fit in comparison to the original SSL theory.

Another limitation of the SSL theory that became obvious when reanalyzing the previous experimental results is its inability to describe how people learn the validities or importance of the different cues. Rakow, Newell, Fayers, and Hersby (2005) showed that the importance participants gave to different cues after learning often corresponded best with the Pearson correlation of the cues with the criterion. However, they did not describe the process by which the subjective importance is learned. Dieckmann and Todd (2004) examined several mechanisms for how people could learn the rank order of cues. One successful and relatively simple rule they studied calls for memorizing a tally of correct and incorrect decisions a cue has made. It would be interesting to incorporate such learning mechanisms into the SSL theory to account for cue-importance learning. However, this is a methodologically challenging enterprise because it requires disentangling the process of learning cue importance from the process of strategy selection learning.

Finally, it should be stressed that although the SSL theory provides a good account of the observed learning process of the experimental studies examined here, there are alternative approaches that do not rely on the assumption of a strategy repertoire that could also explain the observed behavior. Some promising contenders are, among others, the connectionist approach (see, for example, Gluck & Bower, 1988; Sieck & Yates, 2001), the exemplar-based approach (e.g., Juslin, Jones, Olsson, & Winman, 2003; Juslin & Persson, 2002), and the sequential sampling approach (e.g., Busemeyer & Townsend, 1993; Wallsten & Barton, 1982). In this article I did not compare the different approaches with each other (for such studies see Juslin, Karlsson, & Olsson, in press; Persson & Rieskamp, 2007; Rieskamp, 2006a; Rieskamp & Otto, 2006). For such a comparison it is necessary to

design specific experiments in which the different models lead to different predictions.

4.4 Final conclusions

The assumption that people are equipped with a repertoire of different inference strategies has been made in many areas, but the question of how people select different strategies has not been satisfyingly answered. The SSL theory posits that people select strategies on the basis of reinforcement learning. The accuracy of the theory was illustrated by reanalyzing two previous experimental studies. Participants appeared to select their strategies adaptively, such that strategies that performed well were more likely to be selected. In dynamic environments the learning process was slow, so that maladaptive strategies were selected. The SSL theory provides a computational description of how the strategy selection process could take place. To reach adequate conclusions about inference processes the importance of learning should be taken into account.

References

- Beach, L. R., & Mitchell, T. R. (1978). A contingency model for the selection of decision strategies. *Academy of Management Review*, 3, 439–449.
- Betsch, T. (2005). Preference theory: An affect-based approach to recurrent decision making. In T. Betsch & S. Haberstroh (Eds.), *The routines of decision making* (pp. 39–65). Mahwah, NJ: Erlbaum.
- 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. (2006). Adaptive flexibility and maladaptive routines in selecting fast and frugal decision strategies. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 904–918.
- Brown, N. R. (1995). Estimation strategies and the judgment of event frequency. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 21, 1539–1553.
- Brown, N. R., Cui, X., & Gordon, R. D. (2002). Estimating national populations: Cross-cultural differences and availability effects. *Applied Cognitive Psychology*, 16, 811–827.
- Burnham, K. P., & Anderson, D. R. (1998). *Model selection and multimodel inference*. New York: Springer.
- Busemeyer, J. R., & Myung, I. J. (1992). An adaptive approach to human decision making: Learning theory, decision theory, and human performance. *Journal of Experimental Psychology: General*, 121, 177–194.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100, 432–459.
- Christensen-Szalanski, J. J. (1978). Problem solving strategies: A selection mechanism, some implications, and some data. *Organizational Behavior & Human Performance*, 22, 307–323.
- Dieckmann, A., & Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition*, 35, 1801–1813.
- Dieckmann, A., & Todd, P. M. (2004). Simple ways to construct search orders. In K. Forbus, D. Gentner, & R. Regier (Eds.), *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 309–314). Mahwah, NJ: Erlbaum.
- Einhorn, H. J. (1970). The use of nonlinear, noncompensatory models in decision making. *Psychological Bulletin*, 73, 221–230.
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112, 912–931.
- Evans, J. St. B. T., & Over, D. E. (1996). *Rationality and reasoning*. Oxford, England: Erlbaum.
- Fishburn, P. C. (1980). Lexicographic additive differences. *Journal of Mathematical Psychology*, 21, 191–218.
- Garcia-Retamero, R., & Rieskamp, J. (in press). Adaptive mechanisms for treating missing information: A simulation study. *The Psychological Record*.
- 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. (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.
- Gluck, M. A., & Bower, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, 117, 227–247.
- Helverson, B. von, & Rieskamp, J. (2008). The mapping model: A cognitive theory of quantitative estimation. *Journal of Experimental Psychology: General*, 137, 73–96.
- Juslin, P., Jones, S., Olsson, H., & Winman, A. (2003). Cue abstraction and exemplar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 29, 924–941.

- Juslin, P., Karlsson, L., & Olsson, H. (in press). Additive integration of information in multiple cue judgment: A division of labor hypothesis. *Cognition*.
- Juslin, P., & Persson, M. (2002). PROBABILITIES from EXemplars (PROBEX): A “lazy” algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, 26, 563–607.
- 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.
- Mata, R., Schooler, L., & Rieskamp, J. (2007). The aging decision maker: Cognitive aging and the adaptive selection of decision strategies. *Psychology & Aging*, 22, 796–810.
- 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.
- Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior & Human Performance*, 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*. New York: Cambridge University Press.
- Persson, M. & Rieskamp, J. (2007). *Inferences from memory: Strategy- and exemplar-based inference models*. Submitted for publication.
- Rakow, T., Newell, B. R., Fayers, K., & Hersby, M. (2005). Evaluating three criteria for establishing cue-search hierarchies in inferential judgment. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 31, 1088–1104.
- Rapoport, A., & Wallsten, T. S. (1972). Individual decision behavior. *Annual Review of Psychology*, 23, 131–176.
- Rieskamp, J. (2006a). Perspectives of probabilistic inferences: Reinforcement learning and an adaptive network compared. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 1355–1370.
- Rieskamp, J. (2006b). Positive and negative recency effects in retirement savings decisions. *Journal of Experimental Psychology: Applied*, 12, 233–250.
- Rieskamp, J., Busemeyer, J. R., & Laine, T. (2003). How do people learn to allocate resources? Comparing two learning theories. *Journal of Experimental Psychology: Learning, Memory & Cognition*, 29, 1066–1081.
- Rieskamp, J., Busemeyer, J. R., & Mellers, B. (2006). Extending the bounds of rationality: A review of research on preferential choice. *Journal of Economic Literature*, 44, 631–661.
- 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, *Simple heuristics that make us smart*. New York: Oxford University Press.
- Rieskamp, J., & Hoffrage, U. (2008). Inferences under time pressure: How opportunity costs affect strategy selection. *Acta Psychologica*, 127, 258–276.
- Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135, 207–236.
- Selten, R. (1998). Axiomatic characterization of the quadratic scoring rule. *Experimental Economics*, 1, 43–62.
- Sieck, W. R., & Yates, J. F. (2001). Overconfidence effects in category learning: A comparison of connectionist and exemplar memory models. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 27, 1003–1021.
- Siegler, R. S., & Shipley, C. (1995). Variation, selection, and cognitive change. In T. J. Simon & G. S. Halford (Eds.), *Developing cognitive competence: New approaches to process modeling* (pp. 31–76). Hillsdale, NJ: Lawrence Erlbaum.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Smith, V. L., & Walker, J. M. (1993). Monetary rewards and decision cost in experimental economics. *Economic Inquiry*, 31, 245–261.
- Stahl, D. O. (1996). Boundedly rational rule learning in a guessing game. *Games & Economic Behavior*, 16, 303–330.
- Svenson, O. (1979). Process descriptions of decision making. *Organizational Behavior & Human Performance*, 23, 86–112.
- Wallsten, T. S., & Barton, C. (1982). Processing probabilistic multidimensional information for decisions. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 8, 361–384.