# CHAPTER 3

# Taking One Step Back: Simple Ways to Construct Search Orders

## Introduction

In the book *Simple heuristics that make us smart*, Gigerenzer and colleagues (1999) propose several decision making heuristics for predicting which of two objects or options, described by multiple binary cues, scores higher on some quantitative criterion. These heuristics have in common that information search is stopped once one cue is found that discriminates between the alternatives and thus allows an informed decision, leading these heuristics to be termed "one-reason" decision mechanisms. These heuristics differ only in the search rule that determines the order in which information is searched. But where do these search orders come from? Compared to the previous chapters, Chapter 3 therefore takes one step back. So far, the set-up of a heuristic's building blocks and what it takes to be able to apply them has not been given much attention. All questions were about when a certain predefined heuristic, or building block, will be applied. But one might also ask how easily the building blocks can be constructed in the first place. Can simple rules also be at work in the construction of a heuristic's building blocks?

As laid out in earlier chapters, the most prominent of the suggested heuristics, Take the Best (TTB; Gigerenzer & Goldstein, 1996, 1999), searches through cues in the order of their validity. Its performance has been tested on several real-world data sets, ranging from professors' salaries to fish fertility (Czerlinski et al., 1999). Cross-validation comparisons have been made against other more complex strategies, such as multiple linear regression, by training on half of the items in each data set to get estimates of the relevant parameters (e.g., cue order based on validities for TTB, beta-weights for multiple linear regression) and testing on the other half of the data. Despite only using on average a third of the information employed by multiple linear regression, TTB outperformed regression in accuracy when generalizing to the test set (71% vs. 68% correct inferences).

The even simpler one-reason decision making heuristic Minimalist, shortly introduced in the first chapter, was tested in the same way. It differs from TTB only in its search rule: Minimalist searches through cues randomly, and thus requires even less knowledge and precomputation than TTB – all it needs to know are the directions in which the cues point. Again it was surprising that this heuristic performed reasonably close to multiple regression (65% vs. 68%). But the fact that Minimalist lagged behind TTB by a noticeable margin of 6 percentage points indicates that part of the secret of TTB's success lies in its ordered search.

But what is the secret behind ordered search? How can search orders be constructed without assuming full a priori knowledge of ecological cue validities or other environmental statistics?

The remainder of this article will be organized as follows: First, I will argue that research on simple heuristics for cue-based decision making thus far has neglected the process through which people construct an order in which to search for information. In simulations, I will test several, more or less simple cue order learning rules that are proposed as candidates for how such a construction process could look. Whether these rules are descriptive of people's ordering process will be tested in an experimental study. First, on a general level, I will test whether people adapt to the structure of the decision environment both in terms of search, stopping and decision behavior. Then I will explore the search processes in more detail. In particular, I will address questions such as what kind of events per cue affect the order (e.g., discriminations vs. correct decisions), how these events are treated (e.g., size and direction of changes in position within the order), and whether or not one or more tallies per cue are kept and whether a ratio between them is formed (as is necessary for determining validities). These data will provide hints about the kinds of events people pay attention to and thus to the search order construction processes that are at work.

## Experimental evidence for ordered search

From an adaptive point of view, the combination of simplicity and accuracy makes one-reason decision making with ordered search, as in TTB, a plausible candidate for human decision processes. Consequently, TTB has been subjected to several empirical tests. Because TTB explicitly specifies information search as one aspect of decision making, it must be tested in situations in which cue information is not laid out all at once, but has to be searched for one cue at a time, either in the external environment or in memory (Gigerenzer & Todd, 1999).

As summarized in the first chapter, in situations where information must be searched for sequentially in the external environment, particularly when there are direct search costs for accessing each successive cue, considerable use of TTB has been demonstrated (Bröder, 2000, Experiments 3 & 4; Bröder, 2003). This also holds for indirect costs, such as from time pressure (Rieskamp & Hoffrage, 1999), as well as for internal search in memory (Bröder & Schiffer, 2003). The particular search order used has not always been tested separately, but when such an analysis at the level of building blocks has been done, search by cue validity order has received support (Newell & Shanks, 2003, Experiment 2; Newell et al., 2003).

A problem with all these studies has also been pointed out already in the first chapter: None of these experiments tested search rules other than validity ordering. One other very important dimension on which cues can be ordered is discrimination rate, the proportion of all possible decision pairs in which a cue discriminates between the two alternatives. A closer look into the experimental designs of the studies cited above reveals that they all used systematically constructed environments in which discrimination rates of the cues were held constant. Now, when the discrimination rates of cues are all the same, there are not many orders besides validity that make sense. To put it differently, identical discrimination rates make several alternative ordering criteria that combine discrimination rate and validity (such as those listed by Martignon & Hoffrage, 1999) all lead to the same – validity – order. Two examples for such criteria have been introduced in Chapter 1: *success*, the proportion of correct discriminations that a cue makes plus the proportion of correct decisions expected from guessing on the non-discriminating trials (success = $v_i \cdot d_i + 0.5 \cdot (1 - d_i)$, where $v_i$ = validity and $d_i$ = discrimination rate of the cue $i$), and *usefulness,* the portion of correct decisions not including guessing (usefulness = $v_i \cdot d_i$).

Because these criteria collapse to a single order (validity) in the reported experiments, nothing can be said about how validity and discrimination rate may interact to determine the search orders that participants apply. It remains unclear what information participants would base their decisions on when both validity and discrimination rate vary. There are hints that when information is costly, making it sensible to consider both how often a cue will enable a decision (i.e., its discrimination rate) and the validity of those decisions, other criteria such as success that combine the two measures show a better fit to empirical data (e.g., Newell et al., 2004; Läge et al., 2004). But these studies, too, remain silent about how these criteria, or an order based on or reflecting these criteria, could possibly be derived by participants.

In sum, despite accumulating evidence for the use of one-reason decision making heuristics, the basic processes that underlie people's search through information when employing such heuristics remain a mystery. While some clues can be had by considering the size of the overlap or correlations between the search orders people use and various standard search orders (as reported by Newell et al., 2004, and Läge et al., 2004), they do not come close to telling us how cue orders could possibly be constructed.

## Search order construction – the hard way

Although TTB is a very simple heuristic to apply, the set-up of its search rule requires knowledge of the ecological validities of cues. This knowledge is probably not usually available in an explicit precomputed form in the environment, and so must be computed from stored or ongoing experience. Gigerenzer at al. (1999) have been relatively silent about the process by which people might derive validities and other search orders, a shortfall several peers have commented on (e.g., Lipshitz, 2000; Wallin & Gärdenfors, 2000). The criticism that TTB owes much of its strength to rather comprehensive computations necessary for deriving the search order cannot easily be dismissed. Juslin and Persson (2002) pay special attention to the question of how simple and informationally frugal TTB actually is, debating how to take into account the computation of cue validities for deriving the search order. They differentiate two main possibilities on the basis of when cue validities are computed: precomputation during experience, and calculation from memory when needed.

When potential decision criteria are already known at the time objects are encountered in the environment, then relevant validities can be continuously computed and updated with each new object seen. For instance, if you want to predict which soccer team will win the world championships in 2006, you can assess the validity of cues associated with the winning teams in the international matches that take place before the competition. But if it is difficult to predict what decision tasks may arise in the future, this pre-computation of cue validities runs into problems. In that case, at the time of object exposure, all attributes should be treated the same, because any one could later be either a criterion or a cue depending on the decision being made. To use the well-known domain of German cities (Gigerenzer & Goldstein, 1996, 1999; see Table 3.2 on p. 89), the task that one encounters need not be the usual prediction of city populations based on cues such as train connections, but could just as well be which of two cities has an intercity train line based on cues that include city population. To keep track of all possible validities indicating how accurately one attribute can decide about another, the number of precomputed validities would have to be $M^2 - M$, with $M$ denoting to the number of attributes available. In the German cities example, there are 10 attributes (9 cues plus the criterion population size), thus 90 validities would have to be pre-computed. This number rises rapidly with increasing number of attributes. Even ignoring computational complexity, this precomputation approach is not frugal in terms of information storage.

As a second possibility, Juslin and Persson (2002) consider storing all objects (exemplars) encountered along with their attribute values and postponing computation of validities to the point in time when an actual judgment is required. This, however, makes TTB considerably less frugal during its application. The number of pieces of information that would have to be accessed at the time of judgment is the number of attributes times the number of stored objects; in the city example, it is 10 times the number of known objects. With regard to computing validities for each of the $M$ cues, the $N \cdot (N\text{-}1)/2$ possible pairs that can be formed between the $N$ known objects have to be checked to see for each cue if it discriminated, and if it did so correctly. Thus the number of checks to be performed before an order can be determined and a decision can be made is $M \cdot N \cdot (N\text{-}1)/2$, which grows with the square of the number of objects.

Although Juslin and Persson assume worst case scenarios in terms of computational complexity for the sake of their argument, they raise an important point, showing that one of the fundamental questions within the framework of the ABC research group (Gigerenzer et al., 1999) remains open: How can complex, information-hungry processes in determining cue orders be avoided? In other words, can search orders be derived in relatively simple ways?

## Many roads lead (close) to Rome

From what has been said so far, the situation does not look too good for validity either in terms of empirical evidence or psychological feasibility. But what would be the consequence in terms of loss in accuracy if we drop the assumption that cue search follows the validity order? Simulation results can provide an answer. First of all, validity is usually not the best cue ordering that can be achieved. For the German city data set, Martignon and Hoffrage (1999) computed the performance of all possible orderings, assuming one-reason stopping and decision building blocks. The number of possible orders was 362,880 (9! orders of 9 cues). The mean accuracy of the resulting distribution corresponded to the performance expected from Minimalist, 70%, which was considerable above the worst ordering at 62%. Ordering cues by validity led to an accuracy of 74.2%, while the optimal ordering yielded 75.8% accuracy. More than half of all possible cue orders do better than the random order used by Minimalist, and 7,421 (2.0%) do better than the validity order. (Figure 3.1 on p. 90 shows the distribution of all possible cue orders in terms of accuracy and frugality, with TTB's and Minimalist's performance indicated in form of a black and a white star, respectively). It can therefore be concluded that many good orders exist. Moreover, when considering also the costs of a particular cue order in terms of informational demands, cue orders can perform well in several ways, not only by being accurate but also by allowing frugal decision making (such as an order by discrimination rate) or by representing a good compromise between accuracy and frugality (such as success or usefulness order). But how can one of these many reasonably good cue orders be constructed in a psychologically plausible way?

## Search order construction – the simple way

A variety of simple approaches to deriving and continuously updating search orders can be proposed. Indeed, computer scientists have explored a number of self-organizing sequential search heuristics for the purpose of speeding retrieval of items from a sequential list when the relative importance of the items is not known a priori (Bentley & McGeoch, 1985; Rivest, 1976). Consider the case of a set of data records arranged in a list, each of which will be required during a set of retrievals with a particular probability $p_i$. On each retrieval, a key is given (e.g. a record's title) and the list is searched from the front to the end until the desired record, matching that key, is found. The goal is to minimize the mean search time for accessing the records in this list, for which the optimal ordering is in decreasing order of $p_i$. But if these retrieval probabilities are not known ahead of time, how can the list be ordered after each successive retrieval to achieve fast access? Two of the mechanisms proposed to solve this problem are transposition of nearby items and counting of instances of retrieval.[15] The present problem of cue ordering is slightly different from that of the standard sequential list ordering, because cues can fail in ways that retrieved items cannot: A cue may not discriminate (necessitating the search for another cue before a decision can be made), or it may lead to a wrong decision. Still, the mechanisms of transposition and counting will be central to the learning rules proposed below.

The focus is on search order construction processes that are psychologically plausible by being frugal both in terms of information storage and in terms of computation. The decision situation I explore is different from the one assumed by Juslin and Persson (2002) who strongly differentiate between learning of (or about) objects and making decisions. Instead of assuming this unnecessary separation, I will explore a learning-while-doing situation. Certainly there are many occasions akin to Juslin and Persson's situation where individuals have to make decisions based on knowledge they have learned about objects encountered previously and in a different task context. But perhaps more common are tasks that have to be done repeatedly with feedback being obtained after each trial about the adequacy of one's decision. For instance, we can observe on multiple occasions which of two supermarket checkout lines, the one we have chosen or (more likely) another one, is faster, and associate this outcome with cues including the lines' lengths and the ages of their respective cashiers. In such situations, one can learn about the differential usefulness of cues for solving the task via the feedback received over time. It is this case – decisions made repeatedly with the same cues and criterion and the opportunity to learn from outcome feedback – which will now be investigated more closely.

Several explicitly defined cue order learning rules will be considered that are designed to deal with probabilistic inference tasks. Like in the chapters before, the task is forced choice paired comparison, in which a decision maker has to infer which of two objects, each described by a set of binary cues, is "bigger" on a criterion – the task for which TTB was formulated. Thus, in contrast to Juslin and Persson (2002), it is assumed individuals encounter decision situations instead of objects. After an inference has been made, feedback is given about whether a decision was right or wrong. Therefore, the learning algorithm has information about which cues were looked up, whether a cue discriminated, and whether a discriminating cue led to the right or wrong decision. There are different possibilities for taking these pieces of information into account. For example, correct decisions could be counted up for each cue (essentially keeping tallies). Or the information could be used to compute cue validities and discrimination rates based on the cases in which the cue has

[15] A third mechanism is "move-to-front", in which a retrieved record is put at the front of the list, and all other records remain in the same relative order. The search order of the heuristic Take The Last (Gigerenzer & Goldstein, 1996, 1999) rests on this mechanism. There are two variants: One searches for cues in the order of recency of last discrimination, the other by recency of last *correct* discrimination. Both were described in the first chapter of this dissertation, where it was laid out that these rules work well in changing environments. As the focus is on stable environments here, learning rules that are based on the "move-to-front" principle are ignored in the present simulations.

actually been looked up so far. These tallies, validity estimates, etc., would then be used for creating and adjusting the current cue order.

The rules that are proposed differ in the pieces of information they use and how they use them. The learning rules are classified based on their memory requirement – high versus low – and their computational requirements (see Table 3.1). The computational requirements include whether the entire set of cues is completely reordered after each decision or only adjusted locally via swapping of neighboring cue positions, and whether reordering is done on the basis of measures involving division, such as validity, or simple tallying.

*Table 3.1:* Learning rules classified according to their memory requirements and whether they completely reorder the whole set of cues or do local swapping of cue positions.

| High memory load, complete reordering | High memory load, local reordering | Low memory load, local reordering |
|---|---|---|
| Validity: reorders cues based on their current validity<br><br>Tally: reorders cues by number of correct minus incorrect decisions made so far | Tally swap: moves cue up (down) one position if it has made a correct (incorrect) decision if its tally of correct minus incorrect decisions is ≥ (≤) that of next higher (lower) cue | Simple swap: moves cue up one position if it has made a correct decision, and down if it has made an incorrect decision |
| *Variants:*<br>- *reorder based on tally of discriminations so far*<br>- *reorder based on tally of correct decisions only* | *Variants:*<br>- *only upward swapping after correct decisions*<br>- *tally of correct decisions only* | *Variants:*<br>- *moving cues more than one position*<br>- *only upward swapping after correct decisions* |

The *validity rule* is the most demanding of the rules in terms of both memory requirements and computational complexity. It keeps a count of all discriminations made by a cue so far (in all the times that the cue was looked up) and a separate count of all the correct discriminations. Therefore, memory load is comparatively high. The validity of each cue is determined by dividing its current correct discrimination count by its total discrimination count, following Gigerenzer and Goldstein's (1999) suggestion on how people might estimate cue validities (see Chapter 1, p. 19). Based on these values computed after each decision, the rule reorders the whole set of cues from highest to lowest validity.

The *tally rule* only keeps one count per cue, storing the number of correct decisions made by that cue so far minus the number of incorrect decisions. So if a cue discriminates correctly on a given trial, one point is added to its tally. If it leads to an incorrect decision, one point is subtracted from its tally. The tally rule is less demanding both in terms of memory and computation: Only one count is kept, and no division is required.

While the validity and tally rules rely on a counting mechanism, the *simple swap rule* uses the principle of transposition (cf. Bentley & McGeoch, 1985). This rule has no memory of cue performance other than an ordered list of all cues, and just moves a cue up one position in this list whenever it leads to a correct decision, and down if it leads to an incorrect decision. In other words, a correctly deciding cue swaps positions with its nearest neighbor upwards in the cue order, and an incorrectly deciding cue swaps positions with its nearest neighbor downwards.

The *tally swap rule* is a hybrid of the simple swap rule and the tally rule. It keeps a tally of correct minus incorrect discriminations per cue so far (so memory load is medium) but only locally swaps cues: When a cue makes a correct decision and its tally is greater than or equal to that of its upward neighbor, the two cues swap positions. When a cue makes an incorrect decision and its tally is smaller than or equal to that of its downward neighbor, the two cues also swap positions.

As indicated in Table 3.1, many variants of these basic types of learning rules are possible. Here I will focus on these four rules spanning the space of possibilities, and look at how they perform in simulations.

## Simulation study of simple ordering rules

To test the performance of these order learning rules, the German cities data set is used (Gigerenzer & Goldstein, 1996; see Table 3.2), consisting of the 83 highest-population German cities described on 9 cues. The question that will be addressed is, what would happen if a decision-maker does not search for cues in validity order from the beginning, but instead must start with a random order and then construct a search order using feedback received about each decision made? It is assumed that cue directions are known. Furthermore, instead of allowing the decision maker to look up information about all 9 cues in each paired comparison, it is assumed that TTB's stopping and decision rule are used on all decisions constraining what cue values are seen and hence used in order learning. The reason for doing this is that it is more natural to assume that learning happens in the ongoing context of decision making, which does not necessarily involve exhaustive information search. This runs counter the approach taken by Juslin and Persson (2002) who in their worst case scenarios assume exhaustive information search for validity computations. In the approach I follow in this chapter, only the limited information gathered until the first discriminating cue is found can be taken into account.

*Table 3.2:* Cues of the German cities data set (Gigerenzer & Goldstein, 1996) including their ecological validities and discrimination rates.

| Cue | *Ecological validity* | *Discrimination rate* |
|---|---|---|
| National capital? | 1.00 | .02 |
| Exposition site? | .91 | .25 |
| Soccer team? | .87 | .30 |
| Intercity train line? | .78 | .38 |
| State capital? | .77 | .30 |
| License plate? | .75 | .34 |
| University? | .71 | .51 |
| Industrial belt? | .56 | .30 |
| East Germany? | .51 | .27 |

For each rule, 10,000 learning trials were simulated for computing online (cumulative) or amortized performance (Bentley & McGeoch, 1985), defined as the total proportion of correct decisions made so far at any point in the learning process, as well as the contrasting measure of offline performance – how well the current learned cue order would do if it were applied to the entire test set. Each trial started from random initial cue orders, and consisted of 100 decisions between randomly selected decision pairs. For each decision, the current cue order was used to look up cues until a discriminating cue was found, which was used to make the decision. After each decision, the cue order was updated using the particular order-learning rule.

## Results

Figure 3.1 gives an overview of where on average the orders resulting from the learning rules end up in terms of *offline* accuracy and frugality after 100 decisions. This is shown relative to TTB and Minimalist in the distribution of all possible orders. Starting from random orders that on average equal Minimalist's performance, the cue orders produced by the learning rules move both towards higher accuracy and towards higher frugality over time, which can be seen from the fact that the orders learned after 100 trials end up right and below of Minimalist in Figure 3.1.

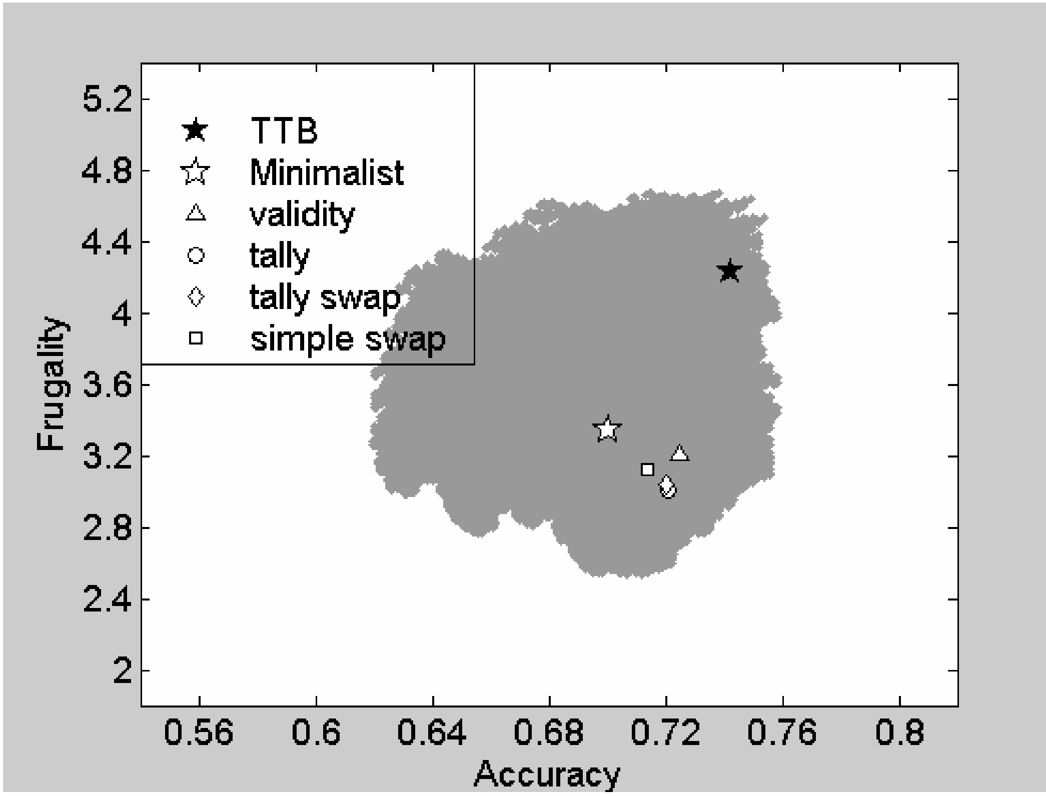


*Figure 3.1:* In grey, all possible search orders for the city size task are plotted in terms of their accuracy and frugality. The stars indicate the performance of TTB and Minimalist. The average offline performance of the cue orders resulting from the learning rules is also marked.

The mean *cumulative* accuracies of the different search order learning rules when used with one-reason decision making are shown on the left part of Figure 3.2. Cumulative accuracies soon rise above that of the Minimalist heuristic (proportion correct = 0.70) which looks up cues in random order and thus serves as a lower benchmark. However, at least throughout the first 100 decisions, cumulative accuracies stay well below the (offline) accuracy that would be achieved by using TTB for all decisions (proportion correct = 0.74), looking up cues in the true order of their ecological validities. The four learning rules all perform on a surprisingly similar level, with less than one percentage point difference in favor of the most demanding rule (i.e., validity) compared to the least (i.e., simple swap; mean proportion correct in 100 decisions: validity learning rule: 0.719; tally: 0.716; tally swap: 0.715; simple swap: 0.711). *Offline* accuracies of the orders resulting after 100 decisions are slightly higher (validity learning rule: 0.725; tally: 0.721; tally swap: 0.720; simple swap: 0.714; see Figure 3.2, right side, as well as Figure 3.1). The more demanding learning rules (i.e., validity and tally) rise above the accuracy of Minimalist more quickly. Besides the mean level of accuracy, an interesting question is whether the learning rules lead to cue orders that converge in terms of accuracy. I therefore computed the standard deviations of the offline accuracy of the resulting orders across the 10.000 runs. Standard deviations decrease for all rules, with a higher decrease for those rules that move closer to accuracy: Starting from 0.026, standard deviations decrease to 0.015 for the validity, tally and tally swap rule, and to 0.022 for the simple swap rule.

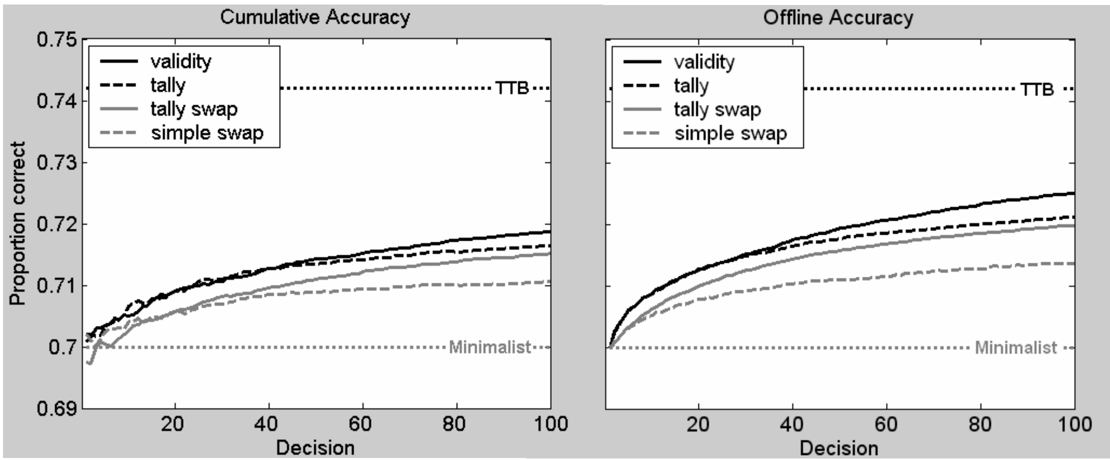


*Figure 3.2:* Mean cumulative accuracy (left) and mean offline accuracy (right) of order learning rules.

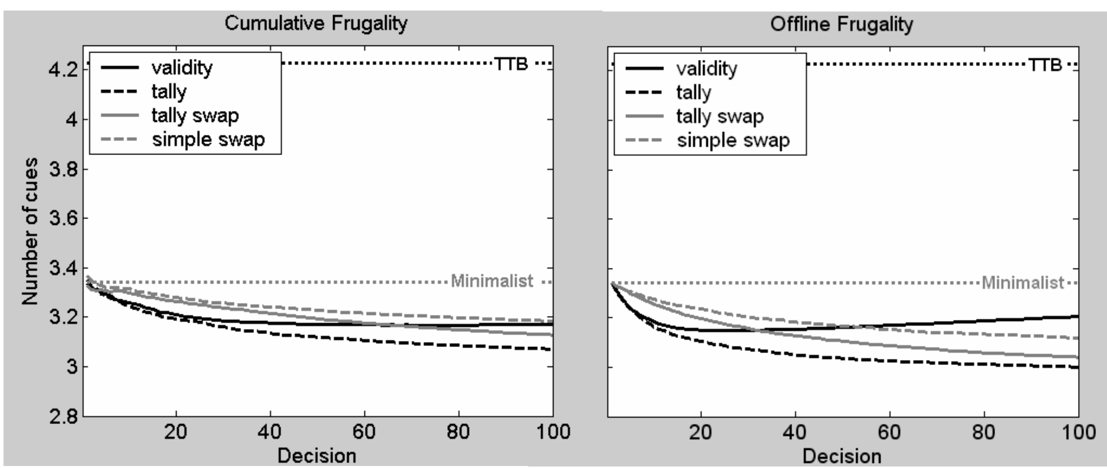


*Figure 3.3:* Mean cumulative frugality (left) and mean offline frugality (right) of order learning rules.

These four learning rules are, however, all more *frugal* than TTB (average number of cues looked up = 4.23), and even more frugal than Minimalist (average number of cues looked up = 3.34), both in terms of online (see Figure 3.3, left side) as well as offline frugality (see Figure 3.3, right side, as well as Figure 3.1). This means that on average, the cue orders produced by these rules look up fewer cues before reaching a decision than do TTB or

Minimalist. Again, there is little difference between the rules (online frugality: mean number of cues looked up in 100 decisions: validity learning rule: 3.17; tally: 3.07; tally swap: 3.13; simple swap: 3.18; offline frugality: mean number of cues looked up by orders resulting after 100 decisions: validity learning rule: 3.21; tally: 3.00; tally swap: 3.04; simple swap: 3.12). The tally rule leads to very frugal cue orders especially quickly. Regarding convergence in terms of frugality, a decrease of standard deviations is again found: Starting from 0.40, standard deviations decrease to 0.24 for the validity and tally rule, to 0.25 for the tally swap rule, and to 0.28 for the simple swap rule.

Consistent with this finding, all of the learning rules lead to cue orders that show positive correlations with the discrimination rate cue order, which is the order that minimizes the number of cues looked up (reaching the following values after 100 decisions: validity learning rule: $r = 0.18$; tally: $r = 0.29$; tally swap: $r = 0.24$; simple swap: $r = 0.18$). This means that cues that more often discriminate between the cities are more likely to end up in the first positions of the order. In contrast, the cue orders resulting from all learning rules but the validity learning rule do not correlate with the validity cue order, and even the correlations of the cue orders resulting from the validity learning rule only reach an average $r = 0.12$ after 100 decisions. Interestingly, in terms of convergence, the standard deviations of these correlations do not decrease to the same extent over the course of the 100 decisions as they do for accuracy and frugality; instead, all stay on a very high level. Let us look at the correlations with discrimination rate as an example: Standard deviations move from 0.36 initially to 0.31 for the validity rule, to 0.33 for the tally and tally swap rule, and to 0.34 for the simple swap rule. This means that although all orders seem to become more similar in terms of accuracy and frugality, this increasing agreement is not achieved by convergence to particular cue orders, because if the cue orders themselves would become more similar to each other, a clear decrease in the standard deviations of the correlations would be observed.

But why would the discrimination rates of cues exert more of a pull on cue order than validity, even when the validity learning rule is applied? Partly, the answer is simple. Remember that the tally rule as well as the swap rules were inspired by self-organizing sequential search heuristics in computer science, designed with the goal to minimize mean search time. To achieve this goal, the rules should be specifically sensitive to discrimination rate. Additionally, there is a more general argument that also holds for the validity rule: Having a low discrimination rate means that a cue has little chance to be used and hence to demonstrate its high validity. Whatever learning rule is used, if such a cue is displaced downward to the lower end of the order by other cues (or starts there in the random initial order), it may never be checked and hence never be able to escape to the higher ranks where it belongs. The problem is that when a decision pair is finally encountered for which that cue would lead to a correct decision, the cue is unlikely to be checked because other, more discriminating although less valid, cues are looked up before it and already bring about a decision. This problem is consequently worse with greater numbers of cues. It is further aggravated by the fact that in the city data set that was used for the simulations, validity and discrimination rate of cues are negatively correlated, such that pulling in the direction of discrimination rate means pulling *against* validity. Thus, because one-reason decision making is intertwined with the learning mechanism and so influences which cues can be learned about, what mainly makes a cue come early in the order is producing a high absolute *surplus* of correct over incorrect decisions – which the tally rule is tracking – and not so much a high *ratio* of correct discriminations to total discriminations regardless of base rates – which validity tracks. I will return to this issue later and relate it to findings on people's processing of frequencies.

In sum, all of the learning rules lead to accuracies between that of the heuristics TTB and Minimalist, but some rules move towards higher accuracy more quickly. The rules are highly frugal, with a (slight) tendency to change cue order in the direction of discrimination rate.

## Discussion

The simpler of the proposed cue order learning rules do not fall far behind a validity learning rule in accuracy. This holds even for the simplest rule, which only requires memory of the last cue order used and moves a cue one position up in that order if it made a correct decision, and down if it made an incorrect decision.

On the other hand, the four rules, even the validity learning rule, stay below TTB's accuracy across a relatively high number of decisions. However, they compensate for this failure by being highly frugal. As it is often necessary to make good decisions without much experience, learning rules may be preferable that quickly lead to orders with good performance. The validity and the tally learning rule quickly move away from Minimalist's performance, towards both higher accuracy and greater frugality, the last aspect being especially true for the tally rule.

The motivation for assessing the above cue order learning rules stems from the necessity of taking into account the set-up costs of a heuristic in addition to its application costs when considering the mechanism's overall simplicity. As seen from Juslin and Persson's (2002) criticism of the standard validity-based search order of TTB, what is easy to apply may not necessarily be so easy to set up. But simple rules can also be at work in the construction of a heuristic's building blocks. Such rules were proposed for the construction of one building block, the search order. It could be shown that these simple learning rules enable a one-reason decision heuristic to perform only slightly worse than if it had full knowledge of cue validities from the very beginning. Giving up the assumption of full a priori knowledge for the slight decrease in accuracy seems like a reasonable bargain: Through the addition of learning rules, one-reason decision heuristics might lose some of their appeal to decision theorists who were surprised by the performance of such simple mechanisms compared to more complex algorithms, but in turn these heuristics can gain psychological plausibility.

Remember that the tally rule as well as the tally swap rule assume memory of the number of all correct minus incorrect decisions made by each cue so far. But this does not make these rules less plausible than a simple swap rule, which, in fact, may not be much simpler, because storing a cue order may be about as demanding as storing a set of tallies. There is considerable evidence that people are actually very good at remembering the frequencies of events. Hasher and Zacks (1984) conclude from a wide range of studies that frequencies are encoded in an automatic way, implying that people are sensitive to this information without intention or special effort. Further, the tally and tally swap rules are comparatively simple, not having to keep track of base rates or perform divisions as does the validity rule. Thus, these rules may represent a good compromise between cost, accuracy and psychological plausibility considerations.

Of course, the crucial test of how psychologically reasonable the tally rule and other rules are as models of human cognition is to see how well they predict people's information search when they have to make cue-based inferences without first knowing validities. Such an experimental test is presented now.

## Experiment

My aim was to find out what information people use, and how they use it, to construct and adjust information search orders in unfamiliar task environments. In almost all previous experimental studies on the use of TTB, participants were encouraged to use cues in order of their validity either directly by informing participants about cue validities or the validity order (Bröder, 2000, Experiments 3 & 4; Bröder, 2003; Bröder & Schiffer, 2003; Newell & Shanks, 2003, Experiment 2; Newell et al., 2003; Rieskamp & Hoffrage, 1999), or indirectly through the presentation of graphs that depicted cue validities (Bröder, 2000, Experiments 1 & 2). Because it is the cue order learning process I am mainly interested in, people were *not* told what the cue validities were in my experiment. Most of the existing experiments on TTB framed the task as choice between differentially profitable shares or stocks from companies that were described on several cues indicative of their profitability (Bröder, 2000, Experiments 3 & 4; Bröder, 2003; Newell & Shanks, 2003; Newell et al., 2003; Rieskamp & Hoffrage, 1999). Because of the potential existence of rather strong initial preferences for certain cues in this relatively familiar domain, I instead created a task most people know very little about: An oil mining situation, in which participants must find out how cues differ in their usefulness for making correct decisions about where to drill for oil. To highlight the importance of searching for the right information in the right order, participants had to pay for information.

## Method

*Stimuli and design*

The study had a 3 x 2 between-participant design. The first factor was the statistical structure of the decision environments: There were 3 different environments, each consisting of 100 decision pairs that could be decided on the basis of 5 cues about the two objects (locations to drill for oil) being compared. In the first environment (in short, *VAL*), cues differed strongly in validity with validities of 0.90, 0.82, 0.73, 0.65, and 0.57, but all had the same discrimination rate of 0.51. In the second environment (*DR*), discrimination rates varied, with values 0.56, 0.49, 0.43, 0.36, and 0.24, while validity was kept constant at 0.75. Finally in a third environment (*VAL*DR*), both discrimination rates and validities varied and were negatively correlated: Validities were 0.57, 0.66, 0.74, 0.83, and 0.91, while the respective discriminations rates followed the opposite order with 0.56, 0.50, 0.43, 0.36, and 0.22.

The environments were constructed in the following way: In principle, with 5 binary cues, 32 distinct cue profiles for objects can be formed. However, making pairs between all these objects yields equal discriminations rates (of 0.51) for all cues. As for two of the environments cues with different discrimination rates were needed, I had to reduce the number of objects with distinct cue profiles to 13 – the maximum number to be found that would yield the chosen discrimination rates. These 13 objects had different criterion values assigned such that the above mentioned cue validities resulted. From these 13 objects in each environment, all 78 possible pairs were constructed. Additionally, in order to produce environments of 100 paired comparisons to allow sufficient time for order learning 22 more pairs were selected randomly from the 78 unique pairs with the condition of leaving validities and discrimination rates largely unaffected. The order of paired comparisons was the same for each participant within one environmental condition. This order was determined by a random process until the following conditions were fulfilled: (a) The cue validities and discrimination rates computed within each block of 20 decisions had to have the same relative order – if there was any – as in the whole data set; (b) at every decision from the 20th onwards, all cues had to have validities, computed based on the paired comparisons encountered so far, no lower than 0.50.

The second factor in the 3 x 2 experimental design was costs for cues, being either high or low relative to gains. Participants received performance-contingent payoff expressed in an artificial currency called “petros”. For each correct decision in the 100 paired comparisons, participants received 2,000 petros (corresponding to 0.20 €). In the high cost condition, they had to pay 300 petros per cue (i.e., relative information costs of 3/20), compared to 100 petros in the low cost condition (i.e., relative information costs of 1/20).

*Hypotheses*

*Hypothesis 1:* Participants' *search order* will move towards the cue order that leads to the highest performance in a given environment, i.e., validity in the first, discrimination rate in the second, and a combination of both (e.g., usefulness or success) in the third environment (but see Hypothesis 2 for a qualification).

*Hypothesis 2:* In the third environment, costs and environmental structure will interact: Discrimination rate will exert a stronger pull on the *search order* under high costs than under low costs, because frugality has a stronger relative impact on net payoff in the high cost condition, while under low costs, accuracy has the stronger relative impact.

*Hypothesis 3:* When costs are relatively high, more *one-reason stopping and deciding* (i.e., stopping cue purchase after the first discriminating cue and deciding accordingly) will be observed than in the low cost condition.

*Hypothesis 4:* The process of search order construction will best be described by *simple learning rules.* In line with the arguments raised in the discussion of the simulation study, I predict that among the simple rules, especially the tally and tally swap rules based on simple frequencies (which excludes the validity learning rule) will achieve a high fit.

*Participants*

120 people (71 female, 49 male; average age 24 years) participated in the study. Most of them were students from the Free University of Berlin. For their participation, participants received performance-contingent payment. The average gain was 13.07 €, and the range was 7.10 to 16.92 €.

*Procedure*

The experiment was completely computer-based. In the instructions participants were asked to imagine they were geologists hired by an oil mining company. Like in the experiments reported in Chapter 2, their basic task was to decide at which of two sites, labeled X and Y, more oil was to be found based on various test results. Besides the framing of the task, there were many crucial differences between this experiment and the ones reported in the second chapter. Therefore, the procedure is again explained in detail.

Five different tests could be conducted, and each had two possible outcomes. Verbal labels were used to represent the dichotomous values, using different labels for each cue. The tests and their respective results were: "Chemical analysis" (for measuring proportion of organic material in ground stone): low vs. high; "Geophones" (for measuring speed of sound waves in ground): slow vs. fast; "Gravimetry" (for measuring strength of changes in gravity relative to surrounding areas): weak vs. strong; "Microscopic analysis" (for measuring pore size in ground stone): small vs. big; "Seismic analysis" (for measuring depth of transitions between stone layers): shallow vs. deep. Although all these tests are indeed used in oil mining (Mobil Oil AG, 1997), they, of course, lead to much more complex results than the dichotomous outcomes I used.

The distinct verbal labels used as cue values were intended to prevent participants from simply counting cases of “yes” (vs. “no”), “+” (vs.”-“) or “1” (vs. “0”) per alternative. The use of verbal labels whose meaning in terms of pro or contra regarding the criterion has to be retrieved seems to be a more realistic way of presenting cue information than numerical values that immediately reveal the cue’s direction and thus represent an already pre-processed type of information. Also, Newell and Shanks (2003) suggested that simple numerical labels for cue values are particularly encouraging to strategies that sum or tally all available information, such as Dawes’ rule (1979), which is a bias I wanted to avoid in the present study.

The tests were first described to participants one at a time, with cue directions revealed by telling participants that at a site with a particular label (e.g. “high” for chemical analysis) more oil is to be found more often than at a site with the opposite label (e.g. “low”). Participants were further told that the tests differed in how reliable they were and in how often they discriminated between sites. In order to facilitate memorization, the stronger adjective (e.g., “big”, “strong”, “fast”, etc.) was consistently used as the positive cue value.

Before the actual decisions started, participants were asked to rank tests according to their presumed usefulness – this was done to be able to check for effects of any preexisting ideas about cue orders. The definition of the word ‘usefulness’ was left open intentionally. To check whether participants had indeed memorized cue directions (which was necessary for making the oil drilling decisions), they then had to indicate for each test separately which of the two different test results most likely pointed to greater oil sources.

After that, the actual decisions started (for a screenshot see Figure 3.4). Participants had to choose between two new oil sites, X and Y, based on the values of test cues that they select to see. Cue values were always revealed pair-wise, that is, simultaneously for both alternatives. At least one test had to be conducted (i.e., one cue had to be selected and revealed). After a test had been conducted, participants could either go on with testing, or decide in favor of one of the sites right away by clicking on the “X” or “Y” button.

The positions at which the tests were displayed at the top of the screen were the same throughout the 100 decisions for one participant, but varied randomly across participants. Test results (cue values) appeared on the lower part of the screen in the order in which the tests were conducted. On which side – left or right – the correct alternative was presented was determined randomly. After a decision had been made by clicking on one of the boxes that represented the alternative drilling sites X and Y, outcome feedback was given: Either a box appeared displaying the word “correct” and the chosen alternative was circled in green, or the box said “wrong” and the chosen alternative was crossed out in red. Furthermore, a cumulative account of participants’ earnings in petros so far was displayed on the screen throughout the decision phase and updated with each cue purchase and correct decision. To decrease differences between high and low cost conditions in terms of final payoff, participants had an account that started at different balances: 10,000 petros in the low, and 30,000 petros in the high cost condition.

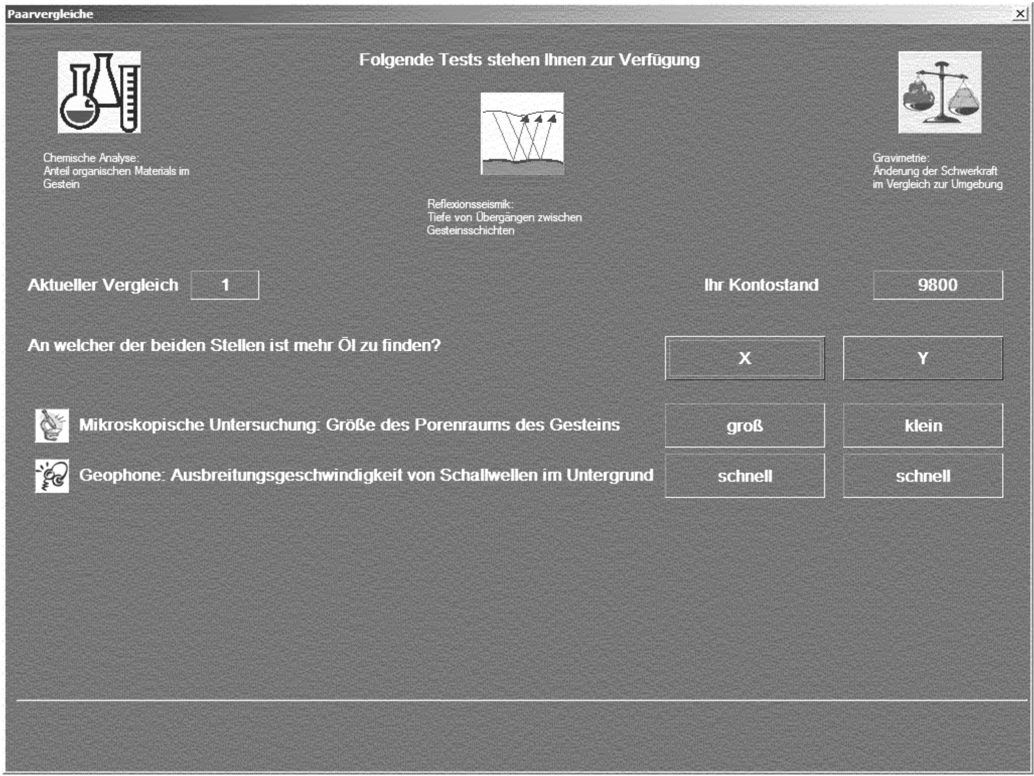


*Figure 3.4:* Screenshot from the experimental program, depicting the task participants faced. This participant, on her first trial, has decided to perform the test "microscopic analysis" (in German: "Mikroskopische Untersuchung") first. Although it discriminates, she performs another test, "geophones", which shows the same the result for both options X and Y. Two hundred "petros" have been withdrawn from her account for conducting these two tests, indicating that she participates in the low cost condition.

After the 100 decisions had been completed, participants were asked to again rank tests according to their 'usefulness'. By doing so, they could increase their gains so far by up to 20,000 petros (i.e., 2 €). They were told about this opportunity for extra reward at the beginning of the experiment in order to motivate cue order learning. The actual payoff was determined by computing the correlation between the participants' indicated rank order and the order that yielded the highest payoff in the particular environment they experienced, and multiplying this correlation by 20,000 petros. Negative payoffs were treated as zero. Finally, there was another test of whether participants had learned cue directions: Just as at the beginning of the experiment, they had to indicate for each test which of the two different test results most likely pointed to greater oil sources. After this, the payoffs resulting from the last cue ranking were added to the gains made during the decision period, converted into €, and paid to participants by the experimenter.

## General results and adherence to TTB's building blocks

First I will report general results and then present data on the building blocks for search, stopping, and deciding, with a special focus on the search rule. These results will allow deriving further predictions about participants' use of simple learning rules. Following this, I will look at the cue ordering process and test whether the suggested simple learning rules are adequate characterizations of what participants actually did.

*General results*

*Cue directions:* I first counted how often participants got a cue value's direction wrong during the pre- and post-experiment quizzes despite having been told this at the beginning. Overall, 120 participants indicated cue directions for five cues two times, resulting in 1,200 answers. Before the decision phase, right after having been informed about the direction into which the cues point, 7 wrong answers were given, and at the end of the experiment, there were 9 wrong answers, corresponding to 99% correctly indicated cue directions overall. Thus, the overwhelming majority of participants knew the directionality of all cues, from the beginning of the experiment until the end.

The average accuracies, frugalities and payoffs participants achieve in the different conditions are summarized in Table 3.3.

*Table 3.3:* Participants' average performance in the different conditions.

| | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|
| | low cost | high cost | low cost | high cost | low cost | high cost |
| Mean percentage correct | *77%* | *74%* | *75%* | *76%* | *71%* | *69%* |
| *(SD)* | *(5.5)* | *(6.2)* | *(7.9)* | *(4.8)* | *(6.2)* | *(6.5)* |
| Mean number of cues seen per trial | 2.69 | 1.92 | 2.79 | 2.28 | 2.78 | 2.28 |
| *(SD)* | *(0.89)* | *(0.48)* | *(0.97)* | *(0.50)* | *(0.83)* | *(0.72)* |
| Mean payoff | 12.79 € | 9.02 € | 12.29 € | 8.41 € | 11.38 € | 6.98 € |
| *(SD)* | *(0.91)* | *(1.36)* | *(1.09)* | *(1.18)* | *(0.85)* | *(1.65)* |

*Frugality*: An analysis of variance with average number of cues bought per trial as dependent variable, and the two between-subjects factors information costs and environment was conducted. In the high cost conditions, fewer cues were bought on average than in the low cost conditions, that is, 2.16 cues on average compared to 2.75, $F(1,114) = 18.62$, $p = .001$, $\eta^2 = 0.14$. The average number of cues bought did not differ between environments (2.30 cues in the first, and 2.53 cues both in the second and third environment), $F(2,114) =$

1.21, $p = .301$, $\eta^2 = 0.02$, nor was there an interaction effect, $F(2,114) = 0.43$, $p = .650$, $\eta^2 = 0.01$.

*Accuracy:* The proportion of correct choices per participant was analyzed in the same way. Despite fewer cues being bought in the high cost conditions, there were no substantial differences in accuracy between cost conditions (73% correct decisions in the high, compared to 75% in the low cost condition), $F(1,114) = 1.66$, $p = .201$, $\eta^2 = 0.01$, and high frugality (as in the high cost conditions) did not incur the cost of low accuracy. However, accuracies differed between environments, $F(2,114) = 11.57$, $p = .001$, $\eta^2 = 0.17$. Post hoc *t*-tests revealed that accuracies of the first environment (where discrimination rate of cues was held constant while validity varied) and the second environment (where, conversely, validity was held constant and discrimination rate varied), with 76% correct decisions on average in both conditions, did not differ, $t(78) = 0.16$, $p = .873$ (two-tailed), $d = 0.04$. Only in the third environment (where both validity and discrimination rate varied and correlated negatively) was a lower accuracy of 70% found, $t(78) = 4.13$, $p = .001$ (two-tailed), $d = 0.92$ in comparison with the first environment, and $t(78) = 4.15$, $p = .001$ (two-tailed), $d = 0.93$ in comparison with the second environment. This result might be due to higher task difficulty in the third environment resulting from the trade-off between validity and discrimination rate. There was no interaction effect, $F(2,114) = 1.29$, $p = .280$, $\eta^2 = 0.02$.

*Payoff:* The payoff achieved by each participant was again analyzed in the same way. Not surprisingly, payoffs achieved in the high cost condition were, with 8.14 € on average, substantially lower than in the low cost condition (12.15 € on average), $F(1,114) = 333.10$, $p = .001$, $\eta^2 = 0.75$. Payoffs also differed between environments, $F(2,114) = 21.32$, $p = .001$, $\eta^2 = 0.27$. Post hoc *t*-tests revealed that payoffs in the first (10.90 €) and in the second environment (10.35 €) did not differ, $t(78) = 1.10$, $p = .273$ (two-tailed), $d = 0.25$. Only in the third environment, a lower payoff of 9.18 € was found, $t(76) = 3.20$, $p = .002$ (two-tailed), $d = 0.72$ in comparison with the first environment, and $t(77) = 2.16$, $p = .034$ (two-tailed), $d = 0.48$ in comparison with the second environment. Thus, between-environment differences in payoff mirror the same differences in accuracy (though payoff is also influenced by frugality of cue use). The was no interaction effect, $F(2,114) = 0.80$, $p = .453$, $\eta^2 = 0.01$.

*To what extent did participants adhere to TTB's building blocks?*

*Search rule*

There were two indicators of the search rule participants actually used by the end of 100 decisions. One was the cue order ranking participants were explicitly asked for at the end of the experiment, and the other was the last search order participants had used in making their decisions. This latter order was derived for each participant in the following way: The cues used on the last, that is, 100th trial were taken in the order in which they were checked. Any missing cues (not checked for the last decision) were added to the end in order of most recent use, so for instance if cue 5 was used on trial 99 but cue 3 had not been used since trial 96, then cue 5 would be followed by cue 3 in the constructed order.

Although the last search order and the indicated rank order are highly correlated (average $r$ = .65), the correlation is not perfect. For most people, the correlation is higher than $r$ = .65 (median $r$ = .80), but there are large individual differences. Participants who do converge to a relatively stable order also tend to report an explicit cue order consistent with this search pattern, whereas this does not hold for some participants who continue to change their search order over and over again until the end of the decision phase. In fact, how much participants changed their search order during the last 20 trials of the decision phase (computed as the number of single position movements necessary to construct one cue order from the previous order) is negatively correlated with the size of the correlation between their last online search order and their indicated search order ($r$ = –.72). I focus here on the search order that participants explicitly indicated after the decision phase. The rationale behind this is that for those people who do converge to a particular order, little information is lost by focusing on the indicated order (as it is highly correlated with the online search order). For those participants who tend to change the cue order throughout the whole decision phase, however, the explicitly indicated order might better represent what they have deduced from the decision phase, whereas their current online order may reflect noisy experimentation. Moreover, unless reported otherwise, the analyses based on the last online search order show the same general pattern.

First, it was checked whether the *initial* explicit ranking participants were asked for was reflected in the final explicit cue order. The correlation between the first indicated and last indicated cue order was on average low (average $r$ = .27). Participants did not even start to search cues in the order they initially indicated, as seen from the correlation between the first explicit ranking and the order in which participants looked up cues for the first time: $r$ = –.05. The correlation between the last indicated cue order and the cue positions on the screen from left to right was also low (average $r$ = .10). It can thus be concluded that neither initial ideas about cue usefulness nor the order in which cues were displayed on the screen had a major impact on the search order participants used.

The basic question posed by Hypothesis 1 is whether people get to a good cue order. The hypothesis is that although people might not learn the best order in a particular environment, their orders will approach the best in each environment. At a minimum, participants should therefore learn cue orders that perform above the level of Minimalist. Is this the case?

The results for the performance of each participants' indicated cue order if applied to all decision pairs they had seen, assuming one-reason stopping and deciding, are summarized in Table 3.4. In the first (equal discrimination rate) environment, the average payoff participants' indicated cue orders would achieve was 13.88 € in the low cost condition and 9.82 € in the high. This is above the mean payoff of all possible cue orders, corresponding to the accuracy of Minimalist, which is 12.96 € in the low cost condition, $t(19) = 4.53$, $p = .001$

(one-tailed), $d$ = 0.89, and 9.26 € in the high cost condition, $t(19) = 2.14$, $p = .023$ (one-tailed), $d$ = 0.48. As discrimination rates of cues did not differ, and therefore different cue orders do not differ in their frugality, this result is explained by higher accuracy of participants' indicated cue orders of on average 79% in the low, and 77% in the high cost condition, whereas Minimalist would achieve an accuracy of 74%.

*Table 3.4:* Average performance of participants' indicated cue orders if applied to all decision pairs in a given environment assuming one-reason stopping and deciding.

| | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|
| | low cost | high cost | low cost | high cost | low cost | high cost |
| Mean percentage correct | 79% | 77% | [79%] | [79%] | 76% | 72% |
| *(SD)* | *(4.5)* | *(5.8)* | *(1.6)* | *(2.1)* | *(5.1)* | *(5.5)* |
| Mean number of cues seen per trial | [1.84 ] | [1.84 ] | 2.06 | 2.00 | 2.27 | 2.08 |
| *(SD)* | *(0.02)* | *(0.03)* | *(0.17)* | *(0.17)* | *(0.23)* | *(0.21)* |
| Mean payoff | 13.88 € | 9.82 € | 13.66 € | 9.77 € | 12.87 € | 8.11 € |
| *(SD)* | *(0.91)* | *(1.16)* | *(0.42)* | *(0.83)* | *(0.85)* | *(0.62)* |

*Note:* Values in brackets refer to numbers that are not expected to be different from a random cue order as validity of cues, or discrimination rate, was held constant in the first and second environment, respectively.

In the second (equal validity) environment, the average payoff of participants' indicated cue orders is 13.66 € in the low cost condition, which is slightly higher than the expected payoff for Minimalist of 13.48 €, $t(19) = 1.87$, $p = .038$ (one-tailed), $d$ = 0.39. In the high cost condition, the average payoff of participants' indicated cue orders is 9.77 €, which is also higher than the 9.25 € expected from Minimalist, $t(19) = 2.79$, $p = .006$ (one-tailed), $d$ = 0.62. In this environment, using an adaptive search order should decrease the number of cues looked up compared to Minimalist. The average number of cues looked up by each participant's indicated cue order if applied to all decision pairs, assuming one-reason stopping and deciding, was 2.06 in the low cost condition and 2.00 in the high cost condition, compared to Minimalist's expected frugality of 2.12 – a small difference. That the payoff of participants' indicated search order still is, at least in the high cost condition, higher than that expected from Minimalist is due to high costs making even small gains in frugality noticeably affect payoff.

In the third environment, both validity and discrimination rate of cues varied. In the low cost condition, participants' indicated cue orders on average achieve a higher payoff than Minimalist, with 12.87 € compared to 12.31 €, $t(19) = 2.93$, $p = .005$ (one-tailed), $d$ = 0.59. This is not the case in the high cost condition, where participants' cue orders achieve

8.11 €, compared to 8.07 € achieved by Minimalist, $t(19) = 0.33$, $p = .371$ (one-tailed), $d = 0.07$. Hypothesis 2 predicted that participants would differ in how they resolve the trade-off inherent in this environment: In the low cost condition, participants should favor cue orders that achieve high accuracy, while in the high cost condition they should favor cue orders that achieve high frugality. In the low cost condition, participants' cue orders are on average more accurate (with 76%) than Minimalist (72%), a gain that has to be paid by looking up more cues (2.27 cues looked up on average, compared to 2.12 for Minimalist). This price is easy to afford in the low cost condition, as the substantial advantage in payoff of participants' cue orders shows. For the high cost condition, however, the pattern does not point in the predicted opposite direction compared to the low cost condition. Rather, the accuracy of 72% and the frugality of 2.08 cues achieved by participants' cue orders are indistinguishable from Minimalist. Maybe, the high costs that had to be paid in this condition have reduced exploration of the environment and as a result, participants were not able to construct an adaptive cue order, which also might have been a more difficult task in this environment compared to the others due to the inherent trade-off between validity and discrimination rate.

Overall, the analysis of the general performance of the indicated cue orders supports the notion that many participants were able to learn an adaptive search order, thus, at least partially, supporting Hypotheses 1 and 2. As a next step, participants' cue orders were correlated with four standard search orders: validity, discrimination rate, usefulness and success. According to Hypothesis 1, participants' cue orders should be adapted to the particular decision environment. Specifically, it was expected that they would correlate with validity cue order in the first, discrimination rate cue order in the second, and success or usefulness cue order in the third environment, in order to respond to the relationship between validity and discrimination rate that existed in each environment.

However, the average rank-order correlations between the participants' indicated cue orders and the standard search orders are quite low (see Table 3.5), and sometimes even negative. Only in the first environment where discrimination rate was kept constant – and high – while validity varied were participants' search orders moderately correlated on average with the ecological validity order (average $r = .36$ in the low, and $r = .30$ in the high cost condition).

To see what the correlations mean in terms of actual overlap between search orders, the orderings of cues were compared starting from first position. Specifically, it was counted how many of the same initial cues were used in two orders (e.g., between orders "2 3 5 1 4" and "2 3 4 1 5" there are two initial overlapping cues). In the first environment under low information costs (where the highest average correlation is found), the average overlap between participants' indicated cue order and the validity search order was 0.80 initial cues. This means that participants' cue orders typically ranked the most valid cue first, but then mostly proceeded to a third- or lower-validity-ranked cue. All other such mean cue overlaps do not exceed 0.60, indicating that participants' indicated cue order often did not even start with the expected cue.

*Table 3.5:* Average correlations of the rank orders participants explicitly indicated at the end of the experiment with four standard search orders.

| | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|
| | low cost | high cost | low cost | high cost | low cost | high cost |
| Validity | 0.36 | 0.30 | | | 0.25 | -0.08 |
| Discrimination rate | | | 0.18 | 0.28 | -0.25 | 0.08 |
| Success | | | | | 0.17 | -0.06 |
| Usefulness | | | | | -0.33 | 0.15 |

While participants' cue orders were not particularly close to the expected objective orders, they were expected to be closer to the *subjective* orders based on their individual learning history. These orders can differ from the objective standard orders, because the base rate of how often participants checked particular cues varied widely, with the most frequently checked cue being looked up on more than 80% of the decisions on average, while the least often looked up cue was seen only in 25% of decision trials on average (see Figure 3.5).

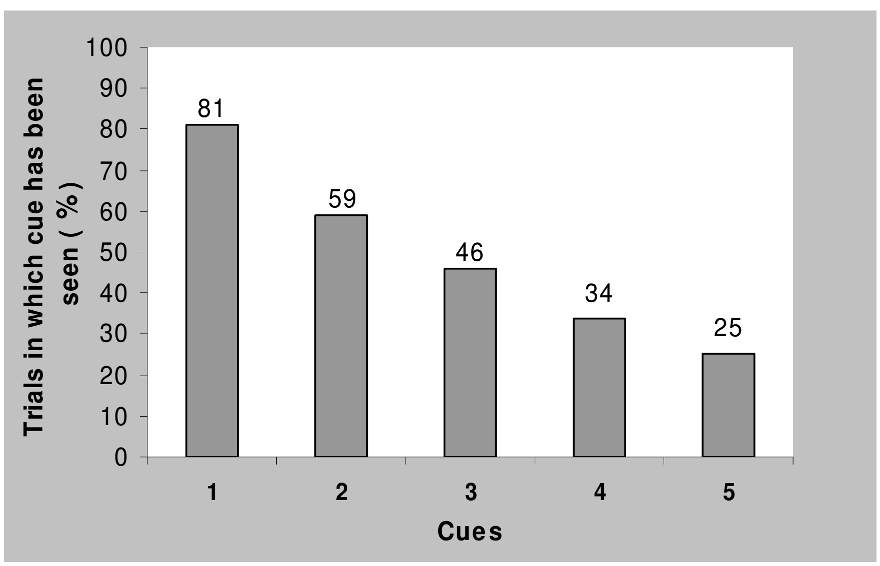


*Figure 3.5:* Percentage of trials in which a particular cue has been looked up, in descending order.

Different base rates for how often each cue was checked result from the learning-while-doing approach of the experimental task – instead of having a separate learning phase in which all cues are presented simultaneously and therefore equally often, participants had to reveal cues sequentially and spend money doing so from the first trial on. Because information search thus became far from exhaustive, certain cues were checked more often than others. Thus, the cues' validities, discrimination rates, success and usefulness that participants actually had experienced were computed by calculating these values across just those cases in which a particular cue had been checked by a participant.[16] When the cue orders based on these subjective measures, computed separately for each participant, were correlated with participants' indicated cue orders, the average rank-order correlations turned out mostly even lower (see Table 3.6), as was also the case for the overlap measure.

Given the surprising failure of participants' cue orders to reflect even these subjective measures, the question arises whether measures like validity might be too complex for this task. Validity is, after all, a conditional probability: the chance that a cue makes a correct decision given that it discriminates. The way the subjective measures were computed also resembles conditional probabilities: Subjective discrimination rate, for example, can be understood as the probability that a cue discriminates given that it had been checked. Taking into account the number of times the cue has been checked makes these subjective measures very similar to ecological validity in terms of computational and memory demands.

*Table 3.6:* Average correlations of the rank orders participants explicitly indicated at the end of the experiment with orders based on *experienced* cue performance measures.

| | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|
| | low cost | high cost | low cost | high cost | low cost | high cost |
| Subjective validity | 0.14 | 0.28 | -0.13 | -0.10 | 0.20 | 0.10 |
| Subjective discrimination rate | -0.19 | -0.25 | 0.08 | 0.07 | -0.27 | -0.04 |
| Subjective success | 0.14 | 0.20 | -0.06 | -0.04 | -0.06 | 0.15 |
| Subjective usefulness | -0.03 | 0.00 | -0.04 | 0.01 | -0.36 | -0.03 |

[16] Specifically, subjective discrimination rates were computed by counting how often a particular cue was looked up and discriminated, and dividing this number by the number of times it had been looked up in total. For usefulness, it was counted how often a particular cue was looked up, discriminated and pointed to the correct alternative; this number was divided by the number of times it had been looked up in total. Success was computed like usefulness with an additionally 0.5 added to the numerator for each time a cue has been looked up and did *not* discriminate. Validity, finally, was computed by dividing the number of times a particular cue was checked, discriminated and pointed to the correct alternative by the number of times that cue was checked and discriminated.

But there are simpler unconditional measures that people could be using instead. To test this possibility, for each participant the number of correct decisions, the number of discriminations, and the number of correct minus the number of wrong decisions they experienced for each cue were simply counted. Again, cues were ordered based on these tallies separately for each participant.

Indeed, the correlations between the participants' search orders and cue orders based on absolute number of correct decisions are higher ($.42 \leq r \leq .66$; see Table 3.7a). Surprisingly, orders based on a tally of mere discriminations, regardless of being right or wrong, also show relatively strong correlations ($.28 \leq r \leq .77$). Only the orders by correct minus wrong tally partly fall behind in terms of the size of the correlations ($-.03 \leq r \leq .55$).

*Table 3.7a:* Average correlations of the rank orders participants explicitly indicated at the end of the experiment with orders based on *experienced* cue performance measures (unconditional ordering criteria).

| | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|
| | low cost | high cost | low cost | high cost | low cost | high cost |
| # correct decisions | 0.66 | 0.66 | 0.54 | 0.57 | 0.42 | 0.67 |
| # correct minus wrong | 0.36 | 0.39 | 0.55 | 0.54 | -0.03 | 0.29 |
| # discriminations | 0.77 | 0.67 | 0.61 | 0.59 | 0.28 | 0.61 |

*Table 3.7b:* Average overlap of the rank orders participants explicitly indicated at the end of the experiment with orders based on *experienced* cue performance measures (unconditional ordering criteria).

| | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|
| | low cost | high cost | low cost | high cost | low cost | high cost |
| # correct decisions | 1.45 | 2.00 | 1.00 | 1.35 | 0.75 | 1.55 |
| # correct minus wrong | 0.75 | 1.00 | 0.70 | 1.15 | 0.20 | 0.65 |
| # discriminations | 2.20 | 1.80 | 1.10 | 1.35 | 0.30 | 1.30 |

*Note:* Cells show the average number of positions participants' indicated cue orders have in common with orders based on different tallies when searching from top to bottom position until one difference is found.

Furthermore, the numbers of initial overlapping positions between the orders are much higher (see Table 3.7b). Thus overall people do seem sensitive to the performance of the cues they have seen, but in a simple way that ignores differences in how *often* they have seen a cue. The differences in fit between groups do not show a consistent pattern. One could for example hypothesize that in the high cost conditions, participants would value the number of mere discriminations a cue makes more than in the low cost condition, where accuracy should be of primary concern. The correlations, however, do not provide any evidence that correct decisions and discriminations are treated differentially depending on the condition, but the ordering process will be examined more closely below.

*Stopping and decision rules*

On the majority of trials, search was stopped immediately after having found one discriminating cue: Across all conditions, the percentage of trials in which participants stopped cue search according to TTB's one-reason rule was 61%. An analysis of variance with percentage of TTB-consistent stopping per participant as the dependent variable and the two between-subjects factors cost and environment was conducted. The proportion of TTB-consistent stopping was with 70% substantially higher in the high compared to the low cost conditions, with 51% (see Figure 3.6), $F(1,114) = 17.19$, $p = .001$, $\eta^2 = 0.13$, but did not differ between environments, (62% in the first, 63% in the second, and 58% in the third environment), $F(2,114) = 0.39$, $p = .679$, $\eta^2 = 0.01$, nor was there an interaction effect, $F(2,114) = 0.07$, $p = .929$, $\eta^2 = 0.00$. At the same time, especially in the low cost condition, participants continued searching after having found one discriminating cue on a considerable proportion of trials (41% in the low cost and 20% in the high cost condition). In contrast, participants rarely stopped before one discriminating cue was found (7% in the low cost and 10% in the high cost condition).

These stopping rule findings nicely replicate those reported by Newell et al. (2003, Experiment 1; see Figure 1.2 in Chapter 1, p. 12) who used six cues and relative cue costs very similar to the high cost condition in the present experiment (1/7, compared to 3/20 in my study). In their paper, they report a higher percentage of TTB-consistent stopping (80%) but this can be attributed to their failure to separate TTB-consistent stopping from stopping *before* a discriminating cue was found. Fortunately, the authors gave me the opportunity to reanalyze their raw data and sort out these cases. Additionally, because it was possible for participants in the Newell et al. (2003) study to guess even before they had checked one cue, a proportion of these cases was counted as potentially being consistent with TTB's stopping rule, assuming that if these guessing participants had been forced to buy at least one cue as in the present study, the participants would then have bought the minimum number of cues required, that is, a single one, and then stopped. On a proportion of those cases equal to the discrimination rate (which was the same for all cues), this behavior would have led to TTB-consistent stopping. This proportion was therefore combined with the proportion of TTB-consistent stopping cases. The resulting overall percentage of trials consistent with TTB in

Newell et al.'s study is then exactly the same as found in the high cost condition of the present experiment: 70%.

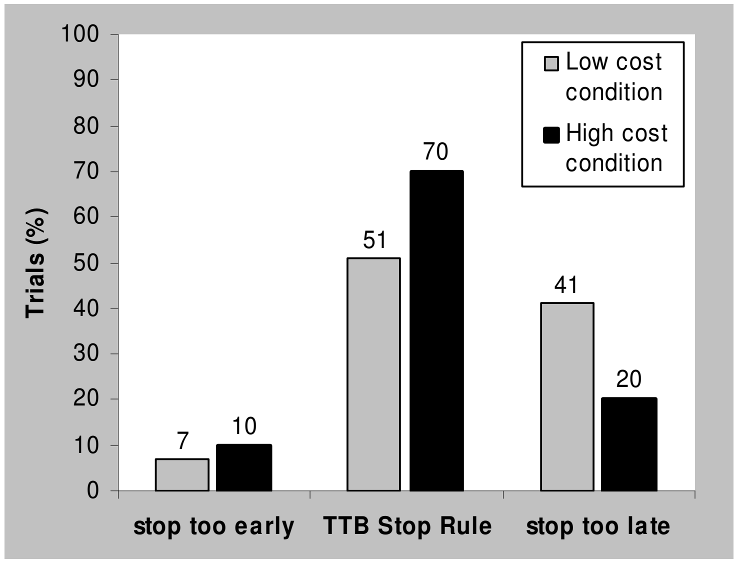


*Figure 3.6:* Adherence to TTB's stopping rule

Does the finding that participants continued searching beyond the first discriminating cue mean that they searched through cues exhaustively in order to apply for instance Dawes's rule or a weighted additive rule? No: on only 25% of the trials where participants searched further than predicted by the TTB stopping rule did they go on to search through all of the cues (i.e., on only 11% of trials overall did participants search exhaustively). But then what were they doing instead? To answer this question, the trials in which search was continued until at least one further discriminating cue was found were analyzed more closely. In half of these cases, the second discriminating cue pointed into the same direction as the first (meaning that in the other half of the cases, it pointed toward the other option). Agreement or disagreement between the two discriminating cues had a strong effect on search continuation: In 83% of all the cases in which the second cue disagreed with the first, search was continued. In strong contrast, participants kept searching in only 14% of the cases where the second discriminating cue pointed into the same direction as the first (see Figure 1.5 in Chapter 1, p. 42). These findings suggest the use of a two-reason stopping rule, or, in short, Take Two, for which evidence was also reported in the second chapter of this dissertation.

How did participants make use of the encountered information to derive at a decision? TTB's decision rule states that one should decide according to the first discriminating cue encountered. Decisions are consistent with this rule on a vast majority of trials (86%; see Figure 3.7 where guesses are defined as cases in which search was stopped before a discriminating cue had been found), again very similar to the results of Newell et al. (2003).

An analysis of variance with percentage of decisions made in accordance with the first discriminating cue per participant as the dependent variable, and the two between-subjects factors cost and environment was conducted. The proportions of decisions made in accordance with the first discriminating cue did not differ between cost conditions (86% in the low vs. 87% in the high cost condition), $F(1,114) = 0.02$, $p = .884$, $\eta^2 = 0.00$, nor between environments (86% in the first, 88% in the second, and 84% in the third environment), $F(2,114) = 0.86$, $p = .428$, $\eta^2 = 0.02$.

How can this high percentage be explained in light of the finding that people frequently search for more information beyond the first discriminating cue? When one-reason stopping is applied, people will most likely also decide in accordance with the single discriminating cue they found. But there are more decisions consistent with one-reason decision making than there are one-reason stopping cases. To see how this happens, one needs to take into account that looking for more information does not necessarily mean that other cues will be encountered that point in the opposite direction – they might just as well confirm the suggestion made by the first discriminating cue, or simply not discriminate. Moreover, even if a second cue is found that points to the opposite alternative, one contradicting cue is probably not enough to outweigh the first discriminating cue.

A closer look at critical trials reveals that in those cases when contradicting information is found that outnumbers the first discriminating cue, one-reason decision making is usually abandoned: Of the few trials (only 248 in total out of 12,000) on which such critical cue value profiles were encountered (e.g., “1 0 0 x x” vs. “0 1 1 x x”), 94% were decided *against* the first discriminating cue. This result also fits with the Take Two rule, mentioned above, indicating that once people apply two-reason stopping, they also decide in favor of the alternative supported by two cues.

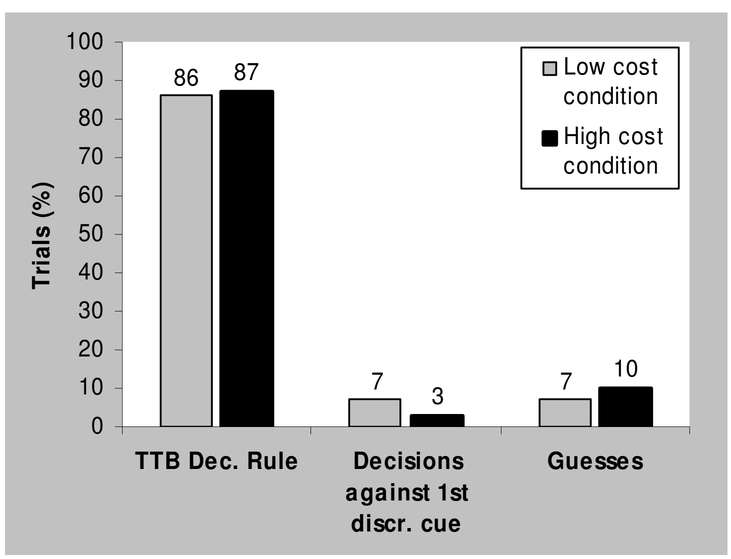


*Figure 3.7:* Adherence to TTB’s decision rule.

## Discussion of the general results

The pattern of results from this experiment suggests that, in accordance with Hypothesis 3, participants in costly decision environments use adaptive stopping rules, usually ending their search after finding one discriminating cue (or sometimes 2 cues supporting the same alternative). Also Hypothesis 1 was partially supported: In most environments, participants' indicated cue orders, if applied to all paired comparisons assuming one-reason stopping and deciding, performed above the level expected from using a random order (Minimalist's performance) in terms of monetary payoff. In the first environment in which cues had constant discrimination rate, this increase in performance was achieved through cue orders with higher accuracy, while in the second, constant validity environment, it was achieved through cue orders with higher frugality (at least in the high cost condition). In the third environment, only partial support for Hypothesis 2 was found. When costs were low, the accuracies of participants' cue orders tended to be higher than Minimalist's while frugalities were lower, resulting in higher monetary payoff than expected from Minimalist. In the high cost condition, however, the predicted opposite pattern was not found: Payoff, as well as both accuracy and frugality, of participants' cue orders were not different from Minimalist.

However, the support for Hypothesis 1 and 2 needs to be qualified: Regardless of condition, the search orders participants use overlap and correlate little on average with standard search orders such as validity, discrimination rate, or success (as suggested by Newell et al., 2004). The only exception is in the first environment: Especially when costs are low and therefore accuracy should be of high concern for participants, their indicated cue orders on average correlate moderately with an order by validity. But average correlations increase when the standard cue orders are replaced by orders following simpler criteria. Participants' cue orders seem to be based on very basic measures such as a simple count of correct decisions, or even mere discriminations, made by a cue. Taking base rates – the frequency with which a cue has been checked – into account when ordering cues, as the orders based on subjective measures do, surprisingly leads to a substantial drop in the average correlations.

In retrospect, it probably should not have come as a surprise that very simple tallying of raw frequencies of cue performance could best match participants' learned cue orders. As discussed earlier, humans are very good at keeping track of frequencies of particular occurrences (Hasher & Zacks, 1984), and this capacity is usually demonstrated in experiments that involve tracking the frequency of many different items (e.g., Flexser & Bower, 1975; Underwood, Zimmerman & Freund, 1971; Zacks, Hasher & Sanft, 1982). Thus, the five cues available in the present experiment were in terms of number of items well within the range for which people are able to give quite accurate frequency judgments. Estes (1976) further argued that people base decisions on these frequencies alone, rather than converting them into base-rate-adjusted probabilities. In a series of experiments, participants first observed outcomes of an imaginary survey about people's preferences for a number of products. Pairs of products were presented, and the product a fictional consumer participating in the survey preferred over the other was indicated. Several different pairs of products were presented (e.g., A vs. B, C vs. D, etc.), with varying frequency, such that probability of winning and frequency of winning of a certain product could be pitted against each other. In the subsequent test phase, critical pairs were formed (e.g., A vs. C, with A having a higher probability of winning, while C has the higher absolute frequency of winning) and participants had to indicate which one was more likely to be preferred by a sample of people from the same population. In this test phase, participants showed a strong tendency to predict that the winner would be the product that had been the winner more frequently in the observation phase, even when it had a lower probability of winning. This result further supports the notion that it might be the number of correct discriminations rather than the conditional measures (i.e., validity and the subjective measures) that participants keep track of and use to determine how good a cue is.

While the correlational results suggest that frequencies of decision events matter for the cue order participants construct, this does not tell us *how* that information is used to build a cue order. Hypothesis 4 therefore remains to be tested. The final cue order alone may be a particularly poor clue to the process that created it, because the learning-by-doing setting leads to unintuitive interactions between the positions of cues and the amount of information that is gathered about them. Therefore, the learning process needs to be examined more closely to find out how participants translate the feedback they receive about the cues during decision making into a cue order. In the following section, I first describe the participants' ordering process in more detail. Then, I test simple cue order learning rules that differ in what information affects cue order and how. For each of the rules, it is computed how well it predicts participants' information search in the present experiment.

## Testing cue order learning rules

*When and how do participants make changes in cue order?*

To get an idea of when and how participants move cues around in their current cue order, changes in the order used from one decision trial to the next will be analyzed. On any given trial $t$, it is assumed that participants have a current cue order, which is inferred as described at the beginning of the results section (pp. 100-101). Then the $M$ cues used in trial $t$+1 are considered to see if they are ordered differently from the first $M$ cues in the current cue order. If so, these changes are related to the cue values and outcome seen on trial $t$. Then, the current assumed cue order for trial $t$+1 is updated, and the analysis proceeds with considering trial $t$+2. The foremost pattern that emerges from this analysis is that cue order usually is *not* changed. On 60% of the trials across all participants, no change in cue position was observed, regardless of the previous decision outcome. Some participants do make many more changes than others, though – the rate of cue order change ranged from 1% to 98% of trials for individuals. Overall, this is congruent with a tendency of participants to converge, more or less quickly, to a particular cue order and then use it for the remaining trials mostly without regard for feedback. I will come back to this point below.

When cues *are* moved, to what extent does their direction of movement follow from the cue's impact on the previous trial? Because of the fact that participants' cue search was in most cases not exhaustive, downward movements are more likely to be observed than upward movements: The first cue, which always had to be looked up, can only be moved down (or stay at its position). Similarly, the last cue can only be moved up (or stay at its position), but as it is not looked up very often, this cannot happen very frequently. This example illustrates that the likelihoods of upward and downward movements are not equal. In order to control for this asymmetry, only cues that were checked at the third position in the search order were looked at. For these, both upward and downward movements are equally possible. When a cue that was looked up at the third position discriminated and indicated a correct decision, it is 1.5 times more likely to be moved up in the order (so it is checked sooner) than to be moved down. In other words, it moved up 28% of the time, stayed in place 54% of the time, and moved down 18% of the time. In contrast, after wrong discriminations a cue is 1.4 times more likely to be moved downwards. When the third cue is checked but does not discriminate, it is also more likely (1.6 times) to be moved downwards.

How far do moving cues travel in the order? The focus is again on cues that were checked at the third position in the search order to control for the unequal likelihoods of upward and downward movements. After correct discriminations, the step size of +1 is the most frequent (besides a step size of 0), at 17%. After non-discriminations, a step size of -1 is most frequently observed (21%) and the same holds for wrong discriminations (21%). Step sizes of +2 and -2 are generally observed rarely. Ignoring the impact a cue had on the previous trial, a travel distance of 2 (regardless of direction) was observed in only 14% of the cases in which a cue had previously been checked at the third position, while the travel distance of 1 was observed in 32% cases overall (with the remaining 54% being instances of no change in position).

These descriptive analyses provide initial hints that people might respond to outcome feedback through adaptive changes to the cue order, that is, moving cues up in the order after correct discriminations, and down after wrong discriminations or after they failed to discriminate. The finding that most often, no change in the search order was observed, regardless of what kind of impact a cue had on the previous trial, potentially speaks against the use of swapping rules and rather supports rules that converge to (relatively) stable orders. Because tally and tally swap rules count up correct decisions or discriminations across all decisions made so far, the relative impact of one single decision decreases over the course of the decision phase. As a consequence, these rules might, consistent with Hypothesis 4, fit behavior better than the simple swap rule. In addition, the relatively high prevalence of step size +1 after correct discriminations and -1 after wrong discriminations could be a hint that tally swap rules rather than complete reordering tally rules might fit behavior well.

*Fit of learning rules*

Finally, it was tested how well different cue order learning or constructing rules could account for participants' ongoing cue search behavior. I tested the same basic types of learning rules as those that were tested in the simulations reported earlier. Motivated by the correlational measures in the general results section, two more variants were added to both the tally rule and the tally swap rule. These variants count correct decisions only and discriminations only, instead of counting correct minus wrong discriminations as in the original tally and tally swap rules. With these four additions there are eight different learning rules.

The fit of the learning rules was computed separately for each participant. For each decision trial, the cue order predicted by each learning rule was compared with the order in which participants looked up cues. After each decision, the current cue order predictions of each learning rule were updated based on the information the participant had encountered in that trial. Unpredictable cases in the first few trials, when no information about a particular cue had yet been gathered (because it had not been looked up), were excluded. That is, the fit was only computed for the cases in which the learning rule made a precise testable prediction about the position of a particular cue.

Two measures of fit were computed for each rule. The first measure is a proportion: Of all cues looked up by a participant on a given trial, how many were seen at exactly the position predicted by the learning rule? The second measure is a distance: Given that a cue was looked up on a particular trial, how far away from its predicted position was it checked? A distance of "2" means that a cue was looked up by a participant two positions away from where the learning rule predicted it would be in the cue order. These distances were averaged across cues, trials and participants.

The average fits achieved by the eight learning rules are reported in Table 3.8 (proportion measure) and Table 3.9 (distance measure). Across all conditions, the set of tally swap learning rules achieve the highest fit in both measures. Within this set, the rule that keeps a tally of correct decisions alone fits best. It correctly predicts the exact position of half of the cues that were looked up (proportion = .50). The mean distance between its predicted positions and where each cue was actually looked up was less than one position (0.87) for this rule. The validity learning rule achieves the lowest fit on both measures (proportion correct: .23; distance: 1.51).

What benchmarks can be used to judge how good the achieved fits really are? One possibility is to predict that the same order will be used on the current trial as was used on the previous trial, for whatever reason. Surprisingly, this "previous order used" benchmark model fits the data slightly better than any of the other rules that were tested (overall fit: proportion measure = .56; distance measure = 0.83). Only in the high cost condition in the first environment were the simple tally rules that apply complete reordering able to achieve an equally high fit.

*Table 3.8:* Proportion of cues that were looked up at exactly the position predicted for it by the respective learning rule, and the corresponding proportion for the proposed benchmark models.

| | **Overall fit** | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|---|
| | | low cost | high cost | low cost | high cost | low cost | high cost |
| Validity | **.23** | .26 | .28 | .20 | .22 | .24 | .19 |
| Tally: | | | | | | | |
| - *correct - wrong* | **.39** | .43 | .53 | .32 | .39 | .29 | .39 |
| - *correct* | **.42** | .47 | .53 | .34 | .42 | .33 | .45 |
| - *discriminations* | **.41** | .50 | .51 | .34 | .41 | .28 | .41 |
| Tally swap: | | | | | | | |
| - *correct - wrong* | **.49** | .58 | .51 | .50 | .50 | .41 | .46 |
| - *correct* | **.50** | .59 | .52 | .51 | .50 | .41 | .49 |
| - *discriminations* | **.49** | .58 | .52 | .47 | .49 | .38 | .48 |
| Simple swap | **.40** | .44 | .41 | .43 | .39 | .35 | .37 |
| Benchmark models: | | | | | | | |
| - *random order* | **.20** | .20 | .20 | .20 | .20 | .20 | .20 |
| - *first order* | **.29** | .34 | .35 | .31 | .21 | .25 | .28 |
| [- *previous order*] | **.56** | .66 | .52 | .61 | .55 | .46 | .52 |

However, this benchmark model and the learning rules rest on very different premises, unfairly favoring the benchmark model: Through resetting the search order to the one most recently used after each decision trial, the "previous order used" model can capture both systematic changes made by participants that rely on processing of feedback as well as unsystematic changes due to errors or fluctuations in the updating rules participants used. In contrast, learning rules that rely on tallies are, by definition, not reset after each decision made. Hence, when a participant makes an error, or changes her updating principle for whatever reason, the learning rules will be unable to predict this change after the next trial. Even worse, unless the participant undoes her unpredicted change, the rule's predictions may fail to match the participant's order for much longer, as the tally values accumulated across many previous trials might predict a different order for several subsequent trials. Thus, pitting the fit of the learning rules against the "previous order used" benchmark is not a fair competition. Out of this consideration, a different version of a benchmark model was formulated that is ignorant of trial-by-trial information. It uses the first search order that is expressed by a participant. This order is reconstructed in the very same way as the last online search order (described on pp. 100-101), only that the *first* complete order that could be constructed was used (e.g., when a participant looked up cues 4 and then 2 on the first trial, cues 2 and 5 on the second, and cues 1, 3 and 5 on the third trial, then her first complete search order, determined after trial 3, would be "1 3 5 2 4"). The performance of this "first order" benchmark is also shown in Tables 3.8 and 3.9. Alternatively, a random model could serve as lower benchmark, with a new random order being generated and applied on each trial. This would lead to an expected proportion of .20 correct predictions, and an expected distance of 1.6 positions.

*Table 3.9:* Average distance of a looked up cue from the position predicted for it by a particular learning rule, and the corresponding results for the proposed benchmark models.

| | **Overall fit** | Environment 1 (VAL) | | Environment 2 (DR) | | Environment 3 (VAL*DR ) | |
|---|---|---|---|---|---|---|---|
| | | low cost | high cost | low cost | high cost | low cost | high cost |
| Validity | **1.51** | 1.39 | 1.38 | 1.51 | 1.58 | 1.48 | 1.69 |
| Tally: | | | | | | | |
| - *correct - wrong* | **1.07** | 0.94 | 0.87 | 1.16 | 1.01 | 1.30 | 1.13 |
| - *correct* | **0.98** | 0.84 | 0.83 | 1.14 | 0.98 | 1.18 | 0.92 |
| - *discriminations* | **1.02** | 0.81 | 0.86 | 1.14 | 1.01 | 1.29 | 0.99 |
| Tally swap: | | | | | | | |
| - *correct - wrong* | **0.89** | 0.63 | 0.93 | 0.85 | 0.91 | 1.07 | 0.92 |
| - *correct* | **0.87** | 0.61 | 0.93 | 0.84 | 0.90 | 1.06 | 0.88 |
| - *discriminations* | **0.89** | 0.63 | 0.93 | 0.89 | 0.92 | 1.11 | 0.90 |
| Simple swap | **0.99** | 0.79 | 1.05 | 0.93 | 1.02 | 1.14 | 1.02 |
| Benchmark models: | | | | | | | |
| - *random order* | **1.60** | 1.60 | 1.60 | 1.60 | 1.60 | 1.60 | 1.60 |
| - *first order* | **1.37** | 1.39 | 1.27 | 1.30 | 1.45 | 1.37 | 1.43 |
| [- *previous order*] | **0.83** | 0.55 | 0.94 | 0.75 | 0.86 | 1.03 | 0.87 |

*Note:* Lower values indicate better fit.

How do the learning rules compare to these lower benchmark models? All of the proposed learning rules achieve a higher fit than expected from a random cue ordering model, although the validity rule's fit is very close to random level, with no difference to be found in some experimental conditions: $t(119) = 2.45$, $p = .008$ (one-tailed), $d = 0.22$ for the proportion achieved by the validity learning rule, for all other learning rules, $t(119) \geq 10.34$, $p = .001$ (one-tailed), $d \geq 0.94$; for the distance measure, $t(119) = 2.26$, $p = .013$ (one-tailed), $d = 0.21$ for the validity learning rule, for all others, $t(119) \geq 12.26$, $p = .001$ (one-tailed), $d \geq 1.12$. But how do the learning rules compare to the much more stable "first order used" benchmark model that assumes that the first search order is used on every subsequent trial? On average, the "first order used" benchmark predicts 29% of cue positions correctly, and the average distance of where a cue has been looked up from its predicted position is 1.37. Thus, all cue order learning rules achieve higher fits than this lower benchmark, $t(119) \geq 4.71$, $p = .001$ (one-tailed), $d \geq 0.43$ for the proportion measure, and $t(119) \geq 5.50$, $p = .001$ (one-tailed), $d \geq 0.50$ for the distance measure, with the exception of the validity rule, which is actually beat by the "first order used" benchmark. Although its fit is worse than that of most of the proposed learning rules, it is surprising that the "first order used" benchmark model can beat the validity rule. A reason for its relatively high fit might simply be that it is a very good description of the search orders applied by a minority of participants who did not bother with changing and updating the cue order, but indeed used the initial search order throughout. This interpretation is supported by the following participant classification results.

Participants were classified uniquely according to the learning rule that predicted the most cue positions correctly for that participant, and that additionally fulfilled the criterion that the proportion of correctly predicted positions is greater than .25.[17] Mirroring the average results across participants, half of the participants fall into the class of tally swap rules (see Table 3.10). Within that set, most participants are classified according to the rule that keeps a tally of correct minus wrong decisions, closely followed by the rule that tracks correct decisions alone. The tally rules that assume complete reordering are best at predicting the cue positions of the search orders used by almost a third of the participants. In harsh contrast, only two participants are classified as following the validity rule, and only one participant as following the simple swap rule. The average proportions of correct cue position predictions achieved for these three participants are also lower than those in the other categories.

It was already suggested above that the first order used, though achieving a relatively low fit on average, might be a good description of the search orders used by participants who do not bother updating their search order. This should be expressed in the form of few people being classified accordingly, but with an on average high fit. Indeed, the search order of ten participants is best predicted by the first search order they expressed, and their average fit, at 66%, the highest of all categories. Another indication is found in the correlation between the first and last online search order: For those classified according to the "first order used" benchmark, this correlation is on average $r = .70$, while it is on average $r = .15$ for all other participants.

Eight participants cannot be classified because the fit of all learning rules fell below the threshold of .25. This indicates that although these participants changed their cue order

---

[17] This threshold was chosen based on the distribution of the proportion of matches expected from a random model. The mean of this distribution is .20, and the standard deviation is .02, so the threshold of .25 is more than two standard deviations away from the mean, and thus expected to be achieved by random choice with a probability less than .01.

throughout the experiment (otherwise the fit of the "first order used" benchmark would be high), these changes were beyond the grasp of any of the learning rules that were tested, thus perhaps indicating random behavior.

*Table 3.10:* Number of participants classified according to each of the learning rule (across all experimental conditions)

| | Number of participants | *Average proportion correctly predicted positions* |
|---|---|---|
| Validity | 2 | *0.41* |
| Tally: | | |
| *- correct - wrong* | 10 | *0.53* |
| *- correct* | 14 | *0.60* |
| *- discriminations* | 13 | *0.50* |
| Tally swap: | | |
| *- correct - wrong* | 25 | *0.58* |
| *- correct* | 23 | *0.62* |
| *- discriminations* | 14 | *0.49* |
| Simple swap | 1 | *0.39* |
| *"First order" benchmark* | 10 | *0.66* |
| *Fit below threshold* | 8 | *0.23* |

*Note:* Participants are classified according to the learning rule that achieves the maximum proportion of correctly predicted positions for each participant individually, under the condition that the proportion (whose means for each rule are shown) exceeds the threshold of .25.

## Discussion of the learning rules' fit

Probably the most striking finding of the analysis of the fit of the cue order learning rules is how clearly the validity learning rule was beat by the other rules, achieving an even lower fit on average than a benchmark model that just assumes that the first complete order is used throughout the decision phase. This result was obtained despite the fact that participants in my experiment could have estimated validity in the online process of decision making in a relatively simple way, contrary to the complex worst case assumptions made by Juslin and Persson (2002). Participants just had to keep two tallies and form a ratio between them. Nevertheless, the validity learning rule failed to predict participants' search behavior.

In line with Hypothesis 4, the learning rules that predict participants' cue order updating process best are those that reorder cues based on simple, unconditional tallies. Among those, the best fits are achieved by rules that refrain from completely reordering the cues. Instead, they just move a cue that has made a correct decision on the most recent trial up one position if its tally is greater or equal to that of the next higher cue. This finding is in line with the research cited earlier on the well developed human capacity for frequency processing (Hasher & Zacks, 1984), and the tendency to sometimes base decisions on raw frequencies of certain outcomes rather than base-rate adjusted probabilities (Estes, 1976). If frequencies are indeed recorded with ease, then this might also explain why the supposedly simplest rule in the competition, simple swap, achieves a relatively low fit. As argued before, storing a cue order, as required by the simple swap rule, may be about as demanding in terms of memory resources as storing a set of tallies, while providing lower performance (as seen from the results of the simulation study). Thus, it is not surprising that participants may have behaved instead in accordance with the more accurate tally and tally swap rules. Additionally, learning rules based on tallies make changes less and less likely over time, thus leading to highly stable orders. In contrast, the simple swap rule does not stabilize the search order, and changes are just as likely at the end of many learning trials as they are at the beginning.

There were no great differences in the fit of the different learning rules across the different environments. This is not too surprising, as tracking the number of correct decisions, or the number of correct minus wrong decisions, will in the long run lead to an adaptive search order in all of the experimental environments, that is, validity in the first, discrimination rate in the second, and a combination of both in the third environment:[18] In the first (equal discrimination rate) environment, more valid cues make more correct decisions *and* less wrong decisions than less valid cues. Thus, tracking these frequencies will favor highly valid cues. In the second (equal validity) environment, cues with high discrimination rate will make more correct decisions than cues with low discrimination rate, and as the ratio of correct to incorrect decisions was the same for all cues, this also holds for the correct minus wrong count. (Here, of course, even just counting discriminations will achieve the same goal.) In the third environment, where discrimination rate and validity vary and correlate negatively, keeping tallies of correct decisions, or correct minus wrong decisions will favor cues that offer a good combination of validity and discrimination rate.[19]

Taken together with the general result that the search orders used by participants mostly achieve a higher payoff than random cue ordering (i.e., Minimalist), the present findings suggest that people adaptively update their cue search order in one-reason decision making, but they do it in simpler ways than prescribed by the validity learning rule.

---

[18] Note though that the correct-minus-wrong tally rule may converge to good orders faster than the correct tally alone, according to the earlier simulations.

[19] The expected number of correct, or correct minus wrong decisions, can easily be computed from the validity and discrimination rate values provided in the description of the experimental environments on p. 95.

Do these results indeed challenge the assumption inherent in TTB that people can construct a search order that follows validity? Or could it be that too little pressure for coming up with the best search order in a given environment can explain why the validity learning rule fits so poorly? Applying the best search order in a given environment would have increased the expected payoff from the 10.90 € achieved by random search to 12.70 €, averaged across all conditions and assuming one-reason stopping and deciding. Some participants might have cared relatively little about continuously improving their search order through careful monitoring of feedback given the already high benefits resulting from adaptive stopping alone, which are more immediately noticeable (i.e., in each trial) without the necessity of slow learning from feedback. Thus, they might have settled with relatively simple cue order updating rules. One could think about ways to make payoffs more contingent on the search rule, by, for example, paying participants more for coming up with a good order. But this would completely change the task in an unnatural direction.

Additionally, the results are restricted to situations where costs for cues exist from the beginning. This might have reduced exploration and the chance to learn more about the cues, including their validity. It remains open what results would be achieved, and what ordering criteria used, when there is a separate cost-free learning phase, and how this would relate to the step-wise process I was interested in. It should be noted, however, that the process of search order construction can hardly be observed once there is a separate learning phase without costs: Any pressure of using a particular order in the learning phase would be removed, because all cues could be looked up in any order without having to pay for it. (This is even more the case if cue presentation is simultaneous.) Nevertheless, the negative results that were found for the validity rule have to be limited to the particular task procedure that was used, and the cue orders resulting from a different kind of learning process – although the process itself would have to remain unspecified – might look different. The possibility that people learn cue validities and construct cue orders based on them under more facilitating circumstances cannot be ruled out. But the low fit of the validity learning rule shows that this is unlikely to happen in the process of a repeated decision making task very similar to the tasks used in the experiments that have demonstrated evidence for TTB use (e.g., Bröder, 2000; Newell et al., 2003), with the only exception that participants in my experiment were not informed about the order of cues by validity.

# General discussion

One goal of the theoretical part of this paper was to argue for taking into account both set-up and application when talking about a decision mechanism's simplicity. To reduce the computational complexity of the set-up of one-reason decision heuristics, simple rules were suggested for the construction of the order in which cues are searched. These rules are inspired by early work in computer science (e.g., Bentley & McGeoch, 1985; Rivest, 1976), where such rules were devised for the problem of repeatedly retrieving items, whose relative importance is not known a priori, from a sequential list. Motivated by the fact that at that time, memory capacity of computers was still a major limiting factor – a constraint that remains important when it comes to *human* cognitive capacities – simple rules that rely on small amounts of stored information were preferred. Simulations have shown that the simple cue order learning rules that were derived enable one-reason decision making to quickly perform better than random search in accuracy and frugality. Among the proposed rules, one that reorders cues after each decision based on their tallies of correct minus wrong decisions performs especially well and very close to a more complex rule based on current validity estimates. At the same time, these simple rules met the second goal of the theoretical part, increasing the psychological plausibility of one-reason decision making by getting rid of the assumption that people know the validity order of cues a priori. The psychological reality of these learning rules themselves was then addressed in an experimental study.

Generally, participants made decisions adaptively in response to various environmental conditions. First of all, this is shown by the analysis of participants' stopping behavior. A high proportion of one-reason stopping (and deciding) was observed, especially when information costs are high. Search was also limited in most of the remaining trials, where participants often searched until two discriminating cues that favor one alternative were found (and then decided accordingly), as predicted by the heuristic Take Two.

Second, the cue orders participants explicitly indicated at the end of the experiment achieved better than random performance in most experimental conditions, even though the decision environments participants encountered had very different statistical characteristics. But at the same time, correlations between participants' cue orders and the standard search orders that would have been best in the different experimental environments (e.g., search by ecological validity) were quite low on average. Rather, participants' cue orders correlated positively with orders based on simple tallies (e.g., of correct decisions only). Thus participants' cue orders, though often beating the random Minimalist strategy, could have done better, lagging behind the respective best performing orders by a relatively wide margin.

The correlational results already suggested that among the cue order learning rules that had been proposed, the rules based on tallies might best depict participants' cue order construction process. Indeed, the fit that these rules achieve supports the notion that they are psychologically plausible descriptions of the cue ordering process. Both types of rules that maintain a record of cues' performance in the form of simple tallies – the tally swap rules and, to a slightly lesser extent, the tally rules that completely reorder the set of cues – predict well the search orders participants use. Within these two classes of rules, the behavior of many participants was best described by those rules that keep tallies of correct decisions alone (or sometimes even discriminations alone). In terms of fit, all of the rules based on tallies clearly beat the simplest rule that had been proposed, that is, the simple swap rule, which just remembers the previous cue order and makes local adjustments based on the outcomes of the previous trial. Keeping tallies may not have demanded too much of participants, and further leads to quite stable orders that after some trials only rarely need to be adjusted. In an even more clear-cut way the validity learning rule was beat by all competitors, which aggravates doubts that people readily learn an order of cues by validity.

What can be concluded from these results? First, the fit of the rules based on tallies supports findings that people are quite good at keeping track of frequencies of events, which could even be an automatic encoding process (Hasher & Zacks, 1984). The assumption of automaticity might also explain why a more complicated ordering criterion, the current validity estimate, achieved a very low fit. Participants might have acquired knowledge about the frequency of correct decisions and the frequency of wrong decisions, necessary to compute the validity estimates, but stored in separate counts. They simply might not have performed the necessary division of the number of correct decisions by the number of discriminations (correct plus wrong decisions) if this required a more conscious form of processing of the potentially automatically acquired frequency information.

A similar explanation may also hold, to a far lesser extent and, as the classification results show, not for all participants, for the rules that keep a tally of correct minus wrong decisions. If participants have stored the number of correct and wrong decisions in separate counts (instead of adding one point to one single tally for a correct decision, and subtracting one point from the same tally for an incorrect decision), this would mean that they did not have available an integrated record. Interestingly, a value account theory of implicit attitude formation has recently been put forward in social psychology (Betsch, Plessner, Schwieren & Gütig, 2001) that assumes that positive and negative responses to an object are summed up in a unitary memory structure. In experiments, the authors presented numerical information about the performance of different shares to their participants. Evidence for the summation model, however, was so far only found for positive events (i.e., for the presentation of shares which exclusively produced gains and no losses), leaving open the question of whether people really keep an integrated account for both positive and negative numerical information about objects.

Alternatively, because wrong decisions in my experiment did not incur a loss of money (besides the money spent on information search), this might have made mistakes less salient, leading people to not actively punish cues for making mistakes. This explanation can be tested in future experiments in which a wrong decision would entail a loss of money just as a correct decision involves a gain (as in the experiments reported in Chapter 2). If the fit of the rules based on correct-minus-wrong tallies were then to increase, it would seemingly be the payoff structure that can explain the results I found in the present study. If not, more fundamental differences in the processing of different events would seem to be implied.

Another conclusion is that in the context of the decision making task I used, participants do not reorder cues by validity. This is in line with the view that validity might be a difficult-to-learn ordering criterion, which calls into question the psychological plausibility argument for validity-based decision making put forward by the simple heuristics approach (e.g., Gigerenzer & Goldstein, 1999; Hertwig & Hoffrage, 2001; Martignon & Hoffrage, 1999). This result was obtained despite the existence of relatively simple ways to estimate validity based on the decisions made so far, in contrast to the scenarios proposed by Juslin and Persson (2002).[20] The decision making task I used was very similar to the tasks participants faced in many experiments that demonstrated the use of TTB, only that – and this is the crucial difference – I did not inform participants about the proper validity order of cues. Of course, this does not rule out the possibility that there are circumstances in which people can and do order cues by validity, or that people acquire validity knowledge from different sources. But where else could people get validity knowledge from?

Knowledge about the validity of certain cues could be the result of evolution, which proceeds slowly, but can create mental mechanisms that require very little computational effort on the part of the individual. Candidate areas to look for such knowledge are, of course, decision problems of evolutionary relevance for which it can be very costly to make mistakes, such as food choice (where high sugar and high fat content, indicating high-energy food, seem to be treated as positive cues, while rotten smell, indicating danger of infectious diseases, is a cue to avoid a certain dish; Curtis & Biran, 2001), avoidance of dangerous animals (where, in rhesus monkeys, even a snake-shaped toy can very easily become associated with a fear response, which then strongly supports the decision not to approach an animal with that form; e.g., Cook & Mineka, 1989, 1990), or deciding who to interact with in social exchange (where past cheating can act as a powerful cue suggesting to avoid interacting with a particular person; Cosmides & Tooby, 1992). Note that in all these examples motivational and emotional responses play a role, possibly making the cues more powerful or even establishing a non-compensatory cue structure – a very adaptive design in high-risk decision domains. However, as these examples already show, it is hard to decide if the cues that are used in these decisions really follow an order by validity, simply because validity can be difficult to determine in some of these cases.

Also, a hierarchy of cues could be the result of social learning or cultural transmission. If we do not know the validity of cues, we can in many cases just look up the necessary data in books or the internet, or directly ask experts. Especially in important and high stake areas,

[20] It should be noted, though, that the validity learning rule did, in the short run, *not* converge on the ecological validity order in the simulations on the German cities data set.

it is likely that someone already has taken the effort to compute validities based on large data sets, e.g., the predictive accuracies of diverse diagnostic cues in medicine, or, to use the example of the experimental task, the validity of various potential indicators of oil deposits. However, for such important decisions, for example when the decision maker will be held accountable for and has to justify the decision, people are prone to gather more information before making a decision and then integrate it (Siegel-Jacobs & Yates, 1996; Tetlock & Boettger, 1989; for a review, see Lerner & Tetlock, 1999). Thus, one-reason decision making is less likely, which in turn reduces the advantages of having a good search order.

Finally, and more in reach of further experimental investigation, individual learning could lead to good estimations of the relative validity of cues. Under some circumstances, people can judge correlations quite well (see Alloy & Tabachnik, 1984, for a review), and ecological validity is a monotonic transform of the Goodman-Kruskal rank correlation ($\gamma = 2v - 1$; see Martignon & Hoffrage, 1999), meaning that both produce the same order. And even if people cannot keep track of the correlations among multiple cues simultaneously, focusing on just the relationship of two variables is more manageable. For example, research on multiple cue probability learning suggests that there might be interference effects when cues are concurrently learned (e.g., Castellan, 1973; Edgell & Hennessey, 1980) that are diminished when cue-criterion relationships are learned one at a time (Brehmer, 1973). Possibly, validities of cues can thus be learned one at a time, and when required by the decision making task, an order of cues based on these validities could be derived ad-hoc.

Along these lines, often our general knowledge about the world might help us focus on certain cue-criterion relationships and thus identify valid cues. Research by García-Retamero, Wallin and Dieckmann (2004) suggests that people make use of causal information about cue-criterion relations as an indicator of highly valid cues. In experiments, the authors found that cues that can be causally linked to the criterion are looked up first, before cues for which such a link is less easily established. Furthermore, participants were more likely to base their decisions on these causally-plausible cues than on others. Causal knowledge might thus reduce the number of correlations to keep track of, through targeting particular cues from an otherwise often wide range of possible cues. This would make the task more similar to a single-cue learning task, in which, as mentioned above, people are better able to learn cue validities than in multiple-cue settings (Brehmer, 1973).

In short, in real-world situations there seem to be several ways to possibly come up with an order of cues by validity. However, in the online process of decision making that was studied in this chapter, no evidence was found that people apply the necessary computations to construct the validity order.

*Conclusion*

Simple cue order learning rules, especially those based on tallies to record ongoing cue performance, enable one-reason decision making heuristics to be quite accurate and very frugal in simulations. However, they do not achieve the accuracy level of TTB, whose search order is based on exhaustive computations of ecological cue validities. When making decisions repeatedly, participants' search order construction process can be quite well described by such tally-based rules, whereas there is no evidence that people reorder cues by validity, even in environments where search by validity would have achieved the highest payoff. This means that so far, empirical evidence for one of TTB's preconditions is still lacking for settings in which an order of cues by validity is not known beforehand. Gigerenzer & Goldstein (1999) wrote: "If people can order cues according to their perceived validities – whether or not this subjective order corresponds to the ecological order – then search can follow this order of cues." (p. 81). It is not clear what exactly the authors have in mind in terms of where these perceived validities come from. But it would be in line with their claim for an external criterion for accuracy (p. 83) to assume that they mean that people's subjective validity orders can be computed *objectively* based on their subjective experience of the cues' successes and failures. It has been shown that in the online process of decision making, people do *not* readily construct search orders based on the validities the cues achieve during the decision making task. It could well be that knowledge about cue validities might often stem from other sources, acquired at different occasions – though it has been shown that these approaches might run into difficulties in terms of computational complexity as well (Juslin & Persson, 2002). But given the low fit the validity learning rule achieved in the experimental study, I believe that the learning rules based on simple tallies describe behavior even under conditions that are more favorable for learning cue validities.