The use of the “Take The Best” Heuristic under different conditions, modeled with ACT-R

Stefani Nellen
Department of Psychology
Ruprecht Karls Universität Heidelberg
Hauptstraße 47-51, 69117 Heidelberg
Stefani.Nellen@urz.uni-heidelberg.de

Abstract

Empirical evidence is accumulating which suggests that the use of the “Take The Best” Heuristic (TTB), one of the “Simple Heuristics that make us smart” (Gigerenzer, Todd & the ABC Research Group) can be induced by factors such as appropriate feedback (Otto & Rieskamp, 2002) or the cost of information search (Bröder, 2000). Using an ACT-R simulation, I demonstrate that these findings can be due to the fact that the utility of heuristics can be learned under some conditions, but not under others: The use of heuristics is conceptualized as being rooted in reaction rather than selection. The theoretical and practical implications of both views are discussed.

Introduction

The investigation of the strategies and heuristics that underlie human reasoning has been a research objective of Cognitive Science for a long time. To identify a set of general reasoning principles would permit us to see more clearly through the jungle of verbal protocols, log-files and program traces produced, e.g. by Solvers of Complex Problems, Students testing Algebra Tutors, or Medical Patients training their scheduling capabilities. Recognizing general patterns in the vast landscape of human cognition would be beneficial for fields such as Education, Software Ergonomics, for the construction of tests and the composition of technical manuals, and many others. However, research which focuses on the mechanics of the mind can run into problems, because it yields very detailed, almost technical theories that are not testable with the less finely grained methods available to psychological research (Anderson & Matessa, 1998). An alternative approach to investigating human cognition is based on the notion that the human mind is shaped by adaptive processes, which means that the mind’s functioning can be related to the structure of the environment(s) it has been functioning in. This view is, e.g., exemplified by Anderson’s rational analysis of memory and strategy-choice (Anderson, 1990, 1998), which has been incorporated into the ACT* theory of skill acquisition (Anderson, 1983) to yield ACT-R (Anderson & Lebiere, 1998). A consequence of this view is that it is legitimate to study the environment, which is observable, in order to make predictions about the behavior of the mind, which is not.

The set of “Simple Heuristics that make us smart” introduced by Gigerenzer, Todd and the ABC Research Group (1999) is another instance of this paradigm. The heuristics presented are not simple and smart per se, but they are “ecologically rational”, i.e. “(they) tap that structure (of different decision environments) to be fast, frugal, accurate and adaptive at the same time” (Todd, Gigerenzer & the ABC research group, 2000, p. 742). In this paper I want to focus on one particular heuristic, the “Take The Best” Heuristic (TTB), because it is among the most intensely investigated of the “Simple Heuristics”. TTB chooses between two alternatives, predicting which of the two will have the higher value with regard to some currently relevant criterion. To do this, the alternatives are compared with regard to their values on cues, or attributes, they both share. The most valid cue is the first to be attended to. “Validity” denotes the de facto correlation of a cue with the criterion of interest (Gigerenzer et al., 1999, p. 46). If that cue discriminates between the two, the alternative it favors is chosen and no further information is needed. If it doesn’t, the second-most valid cue is used, etc. As a consequence, TTB bases its final decision only on the most valid cue that is the first to discriminate between the alternatives: “Take the Best (cue) and ignore the rest”.

TTB can be applied in any situation that involves a choice between two alternatives, provided that these alternatives share some attributes and the decision-maker has at least an intuition about the validity of the cues. Such a situation could, e.g. be the decision to move to one of two cities which can be compared with regard to their wealth, architectural beauty, living expenses etc. The predicted criterion would be “satisfaction with my place of residence”.

Under which conditions is TTB ecologically rational? In addressing this question, Martignon and Hoffrage (1999, p.123) show that “the performance of Take The Best is equivalent to that of a linear model with a noncompensatory set of weights (decaying in the same order as Take The Best’s hierarchy (of cues)). If an environment consists of cues that are noncompensatory when ordered by decreasing validity, then the corresponding weighed linear model cannot outperform the faster and more frugal Take The Best".

When are "Simple Heuristics" used?

Studies investigating the extent in which people use "Simple Heuristics" indicate that, despite their desirable ecological rationality they don't do it quite as often as they could.

Rieskamp & Hoffrage (1999) varied the time pressure in a study on strategy use in a probabilistic inference task (a probabilistic inference task is the kind of decision described above as an example for TTB) and found that under high time pressure, more participants could be classified as using LEX, a variant of TTB, than under low time pressure. (Although it should be noted that the process measures obtained in the analysis of participants' behavior still indicate information search that is far from the "frugality" of TTB; e.g. participants still look up more cues than necessary (p.163)). Bröder (2000, Experiments 2, 3 and 4) confronted his participants with a similar task. He manipulated the conditions "simultaneous vs. successive cue display", i.e. whether participants saw all cue values for the alternatives at once or had too look at them one by one, "feedback (on the quality of the decision) vs. no feedback" and "cost (of looking up the values of additional cues) vs. no cost". Only the combination of successive cue display, feedback and costly information search lead to a noteworthy percentage of participants classified as TTB-users (65% vs. 13% in the other conditions). Finally, Otto & Rieskamp (2002) showed that participants who received feedback that favored decisions in accordance with TTB learned to apply this heuristic- just as participants of another group learned, with the appropriate feedback, to apply a compensatory decision rule[1].

What to make of these findings? One line of reasoning is exemplified by the following quote:

"(...) we consider limited time and limited knowledge as constraints under which people have already developed or learned their smart heuristics. This implies that an individual's repertoire of strategies includes some that take the constraints into account. We do not assume that a trade-off between effort and accuracy or an evaluation of strategies is computed during the decision process. Based on the individual's prior experience of decision making, a particular situation could prompt him or her to use a particular decision strategy." (Rieskamp & Hoffrage, 1999, p.147)

In other words, if the heuristic wasn't used, the conditions have not been sufficient to prompt its use. In principle, however, the heuristic is part of participants' strategic equipment or "adaptive toolbox" (Gigerenzer et al., 1999)

Another view, which I want to put forth in this paper, is to assume that the use of "heuristics" in comparatively novel experimental tasks (such as the ones used in the described studies) is a reaction to experiences made *during* working on that task, rather than the application of something that is already there: the trade-off between effort and accuracy that is associated with certain strategies is not a pre-computed quantity that can be assumed as given, but instead has to be re-assessed as the decision maker enters a novel situation.

I want to illustrate the second view, which might be called "online" or ad-.hoc learning of heuristics (as opposed to "tool application") with an ACT-R simulation in which the model learns to associate TTB with a different utility under the different conditions, resulting in a different extent of its use. Following this demonstration I want to suggest possible ways to test the model presented here, and, finally, discuss its implications for theoretical and, particularly, applied Psychology.

ACT-R Simulation

ACT-R is the offspring of the ACT* theory of skill acquisition (Anderson, 1983) and Rational Analysis (Anderson, 1990). The most recent monograph on the architecture remains "The Atomic Components of thought" (Anderson & Lebiere, 1998). However, a documentation of the (not insubstantial) changes that have since been added to the architecture is available on the World Wide Web.[2] As these changes mainly concern the perceptual/motor components of ACT-R, which are not of immediate relevance to this paper, I will refrain from elaborating them at this time.

In ACT-R, there is a distinction between declarative and procedural knowledge. The availability of the units of declarative knowledge ("chunks") is related to the frequency with which they have been encountered in the environment, and how recently ("Rational Analysis of memory", Anderson & Schooler, 1991). Of more interest here are the units of procedural knowledge: The production rules. Conceptualized as condition/action pairs (which, accordingly, are only applied if their condition is met), their use depends on their expected utility:

$$U = PG - C$$

G is the value of the goal (a little experimented with parameter that is almost always set to the arbitrary value of 20), P is the probability of success and C the cost associated with that production. This cost can be

[1] This seems to contradict Bröders finding of no independent influence of feedback on strategy use; however, an analysis of the stimuli used by Bröder (performed by the author) showed that TTB and a weighed additive strategy arrive almost always at the same decision when faced with them, suggesting that the discriminative value of the feedback may have been insufficient- something Rieskamp and Otto took care to avoid in their study.

[2] at http: // act-r.psy.cmu.edu/papers/403/IntegratedTheory.pdf.

thought of as a “cognitive quantity” (number of problem solving steps that are necessary after the production has fired etc.), but in the simulation here it is used to represent the monetary cost associated with different decision rules. P and G are not fixed parameters, they can be learned from experience[3]. This makes ACT-R particularly appropriate to demonstrate the effect of experiences made during a task on strategy use.

This simulation implements TTB and a weighed additive strategy (WADD) as production rules that can apply in a probabilistic inference task; however, their relative utility isn’t known at the onset of the task but has to be learned. The simulation receives different kinds of feedback, and is faced with different amounts of cost associated with information search.

Task

The task used for the simulation is an abstraction of what might be called the standard probabilistic inference paradigm: A choice has to be made between two alternatives which can be compared on the level of four different cues the validity of which is known. The cues can only be looked up one at a time, in order of their validity (this constraint is adapted from Bröder, 2000). As this simulation is only supposed to illustrate a principle, without reference to a specific task, no particular semantic context was chosen. At any point a decision can be made. A snapshot of this process is depicted in table 1:

Table 1: Snapshot from the decision process in the Abstract probabilistic inference task faced by the simulation. The values of cues 3 and 4 haven’t been looked up yet.

	Alternative A	*Alternative B*
Cue1 validity:.80	yes	no
Cue2 validity:.50	no	yes
Cue3 validity:.30	***	***
Cue -validity:.10	***	***
Choose...	***	***

The model is presented with 120 such decisions. The number and characteristics of these stimuli is adapted from a description in Bröder (2000, Experiment 3, p. 1341). It should be noted that with this set of stimuli TTB has to look up on average one cue in order to make a decision; however, in some cases it has too look up all four of them. WADD always looks up all four cues. The sequence of stimuli presented to the simulation was randomized in each model run to avoid order effects.

[3] In fact, the activation of memory elements and the expected utility of production rules in ACT-R behave quite analogously with regard to learning and decay (Anderson & Lebiere, 1998, ch. 4).

Variation of Feedback and Cost

As shown in table 2, four different combinations of feedback and cost were realized in the simulation: Feedback could be either positive (conditions I and III) or mixed (conditions II and IV).

In order to understand this distinction properly, one should know that, with the stimuli-set used here, TTB and WADD make the same decision in over 90 percent of the cases. Thus “positive feedback” simply means that the model receives the feedback “good choice”, and this applies to TTB as well as WADD. “Mixed feedback”, however, contains a random component, resulting in the feedback “incorrect choice” in approximately a quarter of the decisions. The presumed effect of this manipulation, along with that of the second factor, “cost”, will be explained and demonstrated in the next section.

The term “cost” refers to the monetary cost associated with looking up another cue. The WADD production, which always looks up all four cues, received a cost of 4, while the cost attached to the TTB production was varied, depending on the number of cues TTB had to look up in order to make a decision regarding a specific stimulus.

Table 2: Variations of Cost and Feedback

	Feedback	
	positive	*mixed*
Costly information Search	I	II
Free information search	III	IV

Relation of feedback and cost to the utility and use of production rules

Before reviewing the results of the simulations, it may be useful to summarize the effects that the manipulation of cost and feedback can be assumed to have on the use of TTB and WADD, and why (we can do this because we know the PG-C equation).

Each time a production rule applies, it receives a feedback about the quality of the decision and about the amount of money it has cost. Thus, the nature of the feedback will influence the P Parameter, while the cost (quite obviously) influences the C parameter. As both of these quantities influence the expected utility at the same time, they can not be regarded separately. If two productions are equally successful, the one that is less costly will apply, and if two productions have an equal cost, the more successful will apply. If these quantities are approximately equal, randomness decides initially, until the system settles for one production, which it will continue to apply until the conditions change.

In this particular context, the following patterns in the relative use of TTB and WADD are predicted to appear: If there is an equal cost associated with either of the productions (conditions III and IV), the feedback alone will differentiate between them. Uniformly positive feedback (condition III) will increase the utility of the production rule that is applied first; as there will be no failures, it will be applied more and more frequent. With mixed feedback, however, there is the possibility that one of the productions will receive negative feedback a number of times in a row, so that at a point the expected utility of the competitor is higher, and it is applied instead. On average, The percentage of TTB and WADD should be about equal under that condition (condition IV).
If information search is costly, however, TTB will gradually grow to be the preferred strategy, as it looks up fewer cues than WADD on average. If the feedback is uniformly positive, this preference will emerge quicker, as there will be not setback to its success. If the feedback contains negative elements, the preference should emerge slower and should be less pronounced. Table 3 depicts these basic assumptions.

Table 3: Relative frequency in the use of WADD and TTB under the different conditions

Condition I: Cost/ positive feedback	*Condition II: Cost/ mixed feedback*	*Condition III: No Cost/ positive feedback*	*Condition IV: No Cost/ mixed feedback*
TTB > WADD	TTB >= WADD	TTB >= WADD	TTB = WADD

In the next section, we shall have a closer look at the simulation results, i.e. the relative frequency in the use of TTB and WADD as well as some examples for the process of learning, i.e. the change in the frequency of heuristic use over time.

Results

Use of TTB and WADD under the different conditions

The ACT-R model was run 40 time for each of the four conditions, making 120 decisions during every run. Table 4 shows the percentage of the use of TTB and WADD. The percentage is always relative to the total number of decisions (in this case 40 * 120 = 4800). As table 4 shows, the use of TTB is highest in Condition I and (approximately) equal to that of WADD in condition IV, with the other two conditions hovering in between approximately in concordance

Table 4: Percentage of and WADD use (relative to all decisions) in the different conditions

	Description	TTB	WADD
Condition I	Positive/cost	94	6
Condition II	mixed/cost	73	27
Condition III	Positive/n.c.	69	31
Condition IV	mixed/n.c	47	53

Figures 1 and 2 show the increase in TTB use in condition I and the lack thereof in condition IV, respectively. The stimuli (1st, 2nd,..., 120th decision) are counted on the x-axis and the number of simulated subjects who used TTB and WADD making that particular decision is mapped on the y-axis. In condition I, the model learns to prefer TTB rather quickly (Figure 1), while it keeps switching “confusedly” between TTB and WADD under condition IV (Figure 2).

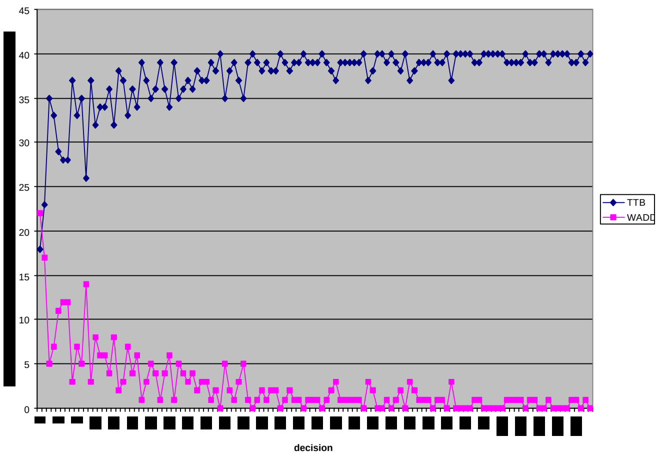

Figure 1: Changes in the use of TTB (triangles) and WADD (squares) under condition I.

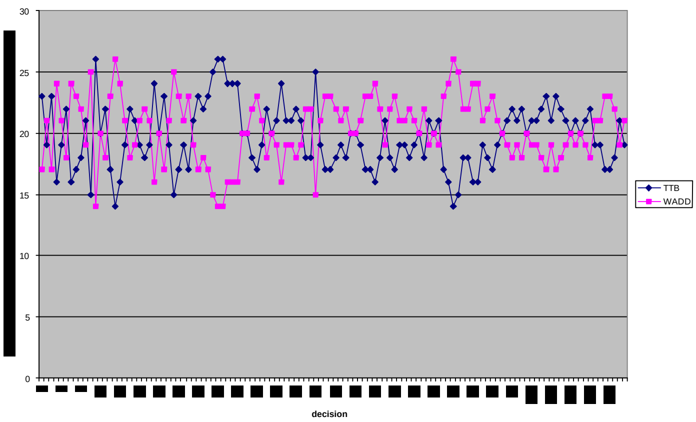

Figure 2: “Changes” in the use of TTB under condition IV.

An additional manipulation: "Confidence" and prior success

These results illustrate the differences in utility learning under various conditions. However, it must be noticed that in all conditions, the percentage TTB was used by the simulation is considerably higher than the one found empirically, e.g. the 13 percent found by Bröder (2000) under conditions of no cost. While this may be due to the simulation being a bit of an oversimplification (e.g.: only two heuristics compete while the number of possible heuristics is much higher, cf. Rieskamp & Hoffrage, 1999, p. 162), another possibility seems to be intriguing. A consequence of the view that the utility of a heuristic like TTB is assessed during the task or ad hoc is that there is either little or no prior experience involving it, or a low probability of success associated with it. In the simulation described above, however, both the WADD and the TTB production were assigned an equal number of initial successes, resulting in them having the same value of P in the beginning. Table 5 shows the percentages of TTB and WADD use under the different conditions, respectively, if TTB is assigned no prior successes, i.e. is assumed to be not associated with any history. Why it is plausible to assume the lack of such a history in the case of TTB but not WADD will be elaborated in the discussion. One interesting result here is a comparison between conditions I and III. If TTB has no prior successes, the combination of its lower cost and positive feedback (I) is needed to make it into the preferred strategy; positive feedback alone (III) is not sufficient, as WADD has a an initial "advantage" before TTB, which, being neither more successful nor less expensive, TTB can't possibly close. As a comparison between conditions II and III shows, the impact of negative feedback is much greater for the strategy with the smaller "prior success" (i.e. TBB). WADD is still applied in 79% of decisions even if it receives as much negative feedback as TTB (IV), while this negative feedback reduces the use of TTB from 94 to 36 %.

Table 5: Percentages of TTB and WADD use if TTB is assigned no prior successes

	Description	TTB	WADD
Condition I	Positive/cost	92	8
Condition II	mixed/cost	36	64
Condition III	Positive/n.c.	6	94
Condition IV	mixed/n.c	21	79

Discussion

In this paper, it has been demonstrated how different conditions can affect the use of heuristics by affecting their utility, which in turn is determined by their cost and success. It has been assumed that these differences in utility are not known in advance but have to be learned first. As a result, a heuristic like TTB is perhaps not applied under conditions in which it would be ecologically rational, simply because the feedback doesn't differentiate between it and other heuristics. An additional condition, such as the cost of information search, would be needed to stress the advantage of TTB over WADD: its parsimony. This assumption disagrees with the notion of an "adaptive toolbox" out of which the appropriate strategies can be selected. Before addressing the theoretical and practical consequences of that view, I want to suggest how to test the model described here.

If participants were presented with decision-situations tin which the application of different heuristics would evoke different responses (see Footnote 1), it would be easy to check whether there is an initial preference for the ecologically rational strategy or whether the strategy that eventually dominates behavior emerges gradually and in response to the manipulations mentioned here.

Another option lies in the investigation of confidence and "prior successes": One could assess participants' confidence in their decisions and whether this relates to the frequency of TTB use, and the speed of acquisition of that strategy in the way the model predicts. The background of these musings was a study by Davis, Lohse & Kottemann (1994), who investigated people's predictions about economical developments. Those participants which were permitted to use more information prior to their decision were more confident regarding its quality- even though the information was redundant. It is at least worth pondering whether "counter-frugal" reasoning, i.e. reasoning that aims at searching out as much information as possible, doesn't have at least as much "ecological credibility" as the "Simple Heuristics". Especially in decisions involving high risk, searching out all information reduces the probability to overlook that one crucial item which – perhaps- will make all the difference. After all, in real life, we often *do* have enough time and money to invest in counter-frugal decision making.

Incidentally, it might be worthwhile to separate the stages of information search and information synthesis, because it is hard to tell whether participants actually use all the information they look up. For this reason it may be appropriate to make limited information search that stops after the "best" cue the diagnostic criterion for the use of TTB. This information search is, e.g. observable by means of successive cue display (as demonstrated by Rieskamp & Hoffrage, 1999), or by means of eye-tracking technology.

The theoretical and practical implications of the "passive" view on strategy as reaction are closely connected and particularly relevant for Instructional Psychology. If heuristics have to be induced by feedback, the resulting "learning" of the heuristics is likely to remain a mere reaction, and transfer of that reaction to other situations is unlikely to occur. This compares unfavorably with the active acquisition of a cognitive skill that is represented explicitly and can be applied in other situations. However, such skills must be taught "the hard way" (see, e.g., Anderson 1987). This calls for the design of instruction methods that do teach decision makers the use of ecologically rational decision rules, e.g. by providing examples and explanations, as well as encouraging the analysis of decision tasks and the criteria for selecting the appropriate heuristic, instead of merely inducing their temporary use. Apart from the challenge of establishing rules that are ecologically rational in the respective situation, this is certainly a worthwhile endeavor. Speed combined with accuracy is a desirable property for any decision rule. Teaching decision-makers the swift selection and correct application of such rules could instruct them to react more appropriately in the real world – and perhaps also increase their confidence in simple rules. That this instruction may be needed is shown by the fact that the participants in the study by Davis et al. (1994), who were so confident in their decisions based on more information, in fact made worse decisions than those who used, or had to use, less.

References

Anderson, J.R. & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Lawrence Erlbaum Associates.

Anderson, J.R. & Matessa, M. (1998). The rational analysis of categorization and the ACT-R architecture. In: Oaksford, M. & Chater, N. (Eds.). *Rational models of Cognition*. New York: Oxford University Press.

Anderson, J.R. & Schooler, L. (1991). Reflections of the environment in memory. *Psychological Science, 2 (6):* 396-408.

Anderson, J.R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Anderson, J.R. (1987). Skill acquisition: Compilation of weak method problem solutions. *Psychological Review 94 (2),* 192 – 210

Anderson, J.R. (1983). *The architecture of cognition.* Hillsdale, NJ: Erlbaum.

Bröder, A. (2000). Assessing the empirical validity of the "Take The Best" heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory and Cognition, 26 (5):* 1332-1346.

Davis, F.D., Lohse, G..L., & Kottemann, J.E. (1994). Harmful effects of seemingly helpful information on forecasts of stock earnings. *Journal of Economic Psychology, 15*: 253-267.

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

Martignon, L. & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In: Gigerenzer, G., Todd, P.M. and the ABC Research Group*: Simple Heuristics that make us smart*. New York: Oxford University Press.

Otto, P. & Rieskamp, J. (2002). Zur Auswahl von Entscheidungsstrategieen. Können einfache Heuristiken erlernt werden? In: van der Meer, E.; Hagendorf, H; Beyer, R.; Krüger, F.; Nuthmann, A. & Schulz, S. (Hrsg.): *43. Kongress der Deutschen Gesellschaft für Psychologie: Programm/Abstracts*. Lengerich: Pabst Science Publishers.

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

Todd, P.M., Gigerenzer, G. & the ABC Research Group (2000). Précis of "Simple Heuristics that make us smart". Behavioral and Brain Sciences 23:5, 727-780.