

How did *Homo Heuristicus* become ecologically rational?

Maria Otworowska (m.otworowska@donders.ru.nl)^a, Marieke Sweers^a, Robin Wellner^a,
Marvin Uhlmann^b, Todd Wareham^c, Iris van Rooij^a

^aRadboud University Nijmegen, Donders Institute for Brain, Cognition, and Behaviour

^bMax Planck Institute for Psycholinguistics

^cDepartment of Computer Science, Memorial University of Newfoundland

Abstract

Gigerenzer and colleagues have proposed the ‘adaptive toolbox of heuristics’ as an account of resource-bounded human decision-making. According to these authors, evolution has endowed such toolboxes with ‘ecological rationality’, defined as the ability to make good quality decisions in their specific environments. Here we explore to what extent the mechanisms of evolution alone can produce ecologically rational toolboxes. We present a formal argument for why evolution is unlikely to produce ecologically rational toolboxes given the astronomically large space of possible toolboxes. The probability of finding one or more ecologically rational toolboxes in this space is negligibly small, even granting an evolutionary time scale of searching for it. We furthermore present artificial evolution simulations results that show that evolution can produce toolboxes of heuristics that are ‘good enough’ to survive, but that those toolboxes are not ecologically rational (in agreement with our formal argument). Our results do not rule out that ontogenetic adaptation processes (development and learning) may yield ecologically rational toolboxes, but it does put into question the idea that phylogenetic processes (evolution) could. We discuss the implications of our findings for future theoretical research on heuristic decision-making.

Keywords: resource-bounded decision making; heuristics; ecological rationality; adaptive toolbox; evolution; computer simulation

Introduction

We make decisions every day, ranging from selecting an outfit or choosing groceries to deciding whom to marry. Even though our decisions aren’t always optimal, they seem to be more often right than wrong in everyday contexts. One prominent account of how we are able to make good quality decisions, despite our bounded resources, is the adaptive toolbox of heuristics account proposed by Gigerenzer and colleagues (Gigerenzer, 2002, 2004, Gigerenzer & Todd, 1999). According to this account, an adaptive toolbox is a collection of specialized cognitive mechanisms—called fast and frugal heuristics—that evolution has built into the human mind for purposes of decision making (Gigerenzer, 2001, Gigerenzer & Sturm, 2012, Gigerenzer & Todd, 1999, p. 30). The heuristics are called ‘fast’ because they can reach decisions with only a few computation steps, and ‘frugal’ because they use little information. Furthermore, the heuristics in the adaptive toolbox are believed to be ‘ecologically rational’ (Gigerenzer, 2002, Gigerenzer & Todd, 1999), i.e. tailored to the contexts in which they are used.

The adaptive toolbox account has had many empirical and explanatory successes in cognitive science (Bergert & Nosofsky, 2007, Bröder, 2000, Dieckmann & Rieskamp, 2007, Goldstein & Gigerenzer, 1999, Pohl, 2006). Yet, the plausibility of the claim that

humans would have evolved adaptive toolboxes of heuristics seems to be so far unexplored. Instead, proponents of the account seem to take the evolutionary plausibility of their cognitive explanation for granted. In this paper we show that the account’s evolutionary plausibility is not self-evident, and even questionable. To see why this is so, we start by considering the notion of ecological rationality as Gigerenzer and colleagues conceptualise it. Next, we explain why evolution is unlikely to produce adaptive toolboxes with the feature of ecological rationality so construed.

Unlike classical notions of rationality that are based on optimality and internal coherence of beliefs and inferences, the adaptive toolbox account defines ecological rationality in terms of the fit between actions and the world. For instance, Gigerenzer & Todd (1999, p. 13) state it as follows: “A heuristic is ecologically rational to the degree that it is adapted to the structure of an environment.” Here, ‘adapted’ refers both to the property of being able to produce actions that fit the environment (i.e., being adapted), and to the process by which the toolbox comes to have that property (i.e., an adaptation process that leads to the property of being adapted to the structure of the environment).

With respect to the fit between heuristics and the environment, Gigerenzer and colleagues claim consistently that this fit (adapted in the property sense) is so good that the quality of decisions is high, and even can outperform optimisation methods (Todd, 2002, Todd & Gigerenzer, 1999, p. 361), at least in those environments to which the heuristics have been adapted (in the process sense). It is because of this good quality that adapted heuristics can be genuinely said to have ecological *rationality*. With respect to the nature of the process of adaptation, two general variants need to be distinguished: phylogenetic adaptation processes (evolution) and ontogenetic adaptation processes (development or learning). Although both types of processes have been claimed to be able to produce adaptive toolboxes that are ecologically rational, here we focus specifically on the (im)plausibility of the idea that a phylogenetic adaptation process would do so.

Clearly, evolution can produce organisms with ecological rationality. By a combination of random variation and selection, organisms can come into existence that have decision tendencies that are particularly tuned to particular environments. However, it is highly implausible, that organisms (especially humans) would come to have such high degrees of ‘fitness’ if their decisions were based on toolboxes of heuristics and evolution was to set the parameters of these toolboxes directly. The reason is that toolboxes of heuristics have

an enormous amount of degrees of freedom: A toolbox can vary in terms of the number of heuristics it contains, and each heuristic can vary in terms of both the possible environmental cues to which it responds and the different possible actions it can perform. Given that the number of possible cue-heuristic-action mappings grows exponentially in these parameters, the number of distinct possible toolboxes does as well.

Given these considerations, what are the odds of evolution producing toolboxes that are ecologically rational? This depends on how many toolboxes in the vast space of possible toolboxes are ecologically rational. As we will show, the vast majority of possible toolboxes aren't ecologically rational. Even though the mechanisms of natural selection are not random, the only evolutionary mechanisms that can produce different toolboxes—such as mutation and crossover—are random. This means that the chance of creating, and subsequently selecting, ecologically rational toolboxes is so nanoscopically small that even on an evolutionary time scale it is extremely improbable that evolution would yield ecologically rational toolboxes. In this paper, we elaborate on this argument both formally and using computer simulations.

The remainder of this paper is organized as follows. We present a formalization of the notion of an adaptive toolbox, to be used both in our formal argument and our computer simulations. Next, we put forth a formal argument for the implausibility of the idea that evolution could produce ecologically rational toolboxes based on illustrative numerical estimates for even small toolboxes. We then describe the setup of an artificial evolution environment that we use to empirically validate our argument. We present results of simulations for three different setups, each demonstrating that even though evolution can produce toolboxes that are 'good enough' to survive, these toolboxes do not display any notable ecological rationality. We close by discussing the broader implications of our findings for research into resource-bounded decision making.

Formalizing the Adaptive Toolbox

In this section we will present a formalization of the adaptive toolbox account, which involves formalizing components of the adaptive toolbox (heuristics with a selector) as well as its environment. We represent each of the components as a fast and frugal tree (see Figure 1). Each internal node in such a tree stands for a boolean function; a tree evaluates only a limited set of statements (cues; which can be either true or false) and a particular action is triggered by a particular sequence of cues progressing from the root-node to the leaf representing that action.

Environment

The environment consists of a set of *events* (environmental cues) $E = \{e_1, e_2, \dots, e_n\}$, every event can be either true or false. A truth assignment for each event is called a *situation* s . That is, a function s assigns truth values to each event in E , $s : E \rightarrow \{T, F\}$. We denote the set of all possible situations by $S = \{T, F\}^n$, where S is the set of all possible n -length vectors

of truth-values. For every situation there is a certain favored *action* a to perform, where a is an element of the set of all possible actions $A = \{a_1, a_2, \dots, a_m\}$. A function $D : S \rightarrow A$ maps each situation $s \in S$ to an action $a \in A$.

Heuristics

Each heuristic in the toolbox is represented as a *fast and frugal tree* (Gigerenzer & Gaissmaier, 2011, Martignon *et al.*, 2003), a chain of cues with associated actions. Each cue is a boolean function, evaluating whether an event $e \in E$ is true in a given situation, $c(e, s)$. When executing a heuristic, the tree is traversed starting at the top. Step by step the cue functions are passed, checking whether the cue holds. If the cue $c(e, s)$ evaluates to true for event e is in situation s , then the action a associated to that cue c is executed. If the cue is false the next cue is evaluated until the bottom cue is reached. If this last cue is false, the last action in the tree is performed.




Figure 1: A single heuristic represented as a fast and frugal tree. The tree contains cues $C = \{c_1, c_2\}$ and associated actions $\{a_1, a_2, a_3\}$. Each cue $c \in C$ is a simple boolean function which evaluates whether an event $e_j \in E$ is true or false, depending on the situation $c(e_j, s_k)$. If the cue function returns 'true', the tree traversal stops and the action associated with the cue is executed; otherwise the next cue function is executed. For example, if c_1 is false, but c_2 is true, then the action a_2 will be executed.

Selector

A selector determines which heuristic to use in a given situation. We represent the selector as a fast and frugal tree as well¹; the internal nodes are cues associated with heuristics (see Figure 2). A heuristic is executed in the case a cue is evaluated to be true.

Mathematical analysis

In this section we present a formal argument for the implausibility of generating the ecologically rational adaptive toolboxes by means of evolutionary processes alone. The argument is composed of three parts: search space argument, probability argument and time argument.

¹Hypotheses about the exact nature of the selector mechanism haven't been developed to the same extent as hypotheses about the structure of individual heuristics. Nevertheless, the common idea seems to be that the selector, like the heuristics, is fast and frugal. For our purposes, and without loss of generality, we assume that the selector can be modelled by a fast and frugal tree as well.




Figure 2: The adaptive toolbox selector and heuristics as fast and frugal trees. The selector is represented by the orange nodes. The tree is traversed from left to right (selecting a heuristic) and from top to bottom (executing a heuristic). For instance, let’s assume a situation such that c_4 (“it is sunny outside”) = F , $\neg c_2$ (“I have not read any book in a while”) = T , and c_5 (“my favorite book is on the shelf”) = T ; in such a case the action a_2 = “read the book” will be executed. Note that when the last cue of the selector (c_3) returns false, the first heuristic is executed by default.

Part 1: Search space and location-sensitivity

Let’s assume a simple environment (10 events, 50 actions).² For the purpose of the analysis we use the simplification that environments are structured such that at least one adaptive toolbox would be able to act perfectly in it. Then there are $2^{10} = 1024$ situations an individual may encounter during its lifetime (see section Environment). Further, let’s assume a simple toolbox of a size $12 = 3$ (number of selector cues) + $(3$ (number of heuristics) $\times 3$ (number of cue/action pairs in each heuristic)). The number of all possible different toolboxes is $10^{12}(\text{cues}) \times 50^9(\text{actions}) = 10^{27}$. Let’s consider a toolbox to be ecologically rational if it performs actions which are more often right than wrong. Given that we define the fitness score as the proportion of the number of situations in which a toolbox executes a correct action to the total number of all possible situations, the fitness is in a range 0 to 1 inclusive, and a score of ≥ 0.5 indicates ecological rationality.

Table 1a represents a toolbox of size 12. We set the probability of a given cue being true or false to 0.5. That means that for the first cue of the selector (S1 in the Table 1a) there is a 50% chance that it will be true (and the first heuristic will be executed) and 50% chance that it will be false (and the next selector (S2) cue will be evaluated). We can now estimate the degree to which cues and actions contribute to the toolbox’s fitness as a function of their location in the toolbox.

²Here, 50 actions may seem like a lot, but taking into account the number of different things one can do e.g. with any given object (grasp it, throw it, squeeze, cut it, etc.) it is actually a moderate estimate.

(a)	S1		S2		S3	
	H1:C1	H1:A1	H2:C1	H2:A1	H3:C1	H3:A1
	H1:C2	H1:A2	H2:C2	H2:A2	H3:C2	H3:A2
	H1:C3	H1:A3	H2:C3	H2:A3	H3:C3	H3:A3

(b)	50%		25%		12.5%	
	25%	25%	12.5%	12.5%	6.25%	6.25%
	12.5%	12.5%	6.25%	6.25%	3.125%	3.125%
	6.25%	6.25%	3.125%	3.125%	1.6%	1.6%

Table 1: (a) A schematic representation of a toolbox of size 12. In this toolbox, S1 is the first selector cue, H1:X is the first heuristic, H1:C1 is the first heuristic cue and H1:A1 is the first action in the first heuristic. (b) A representation of contribution of cues and actions to fitness depending on their locations in a toolbox. The blue color indicates the minimal requirement for an ecologically rational toolbox.

If the first selector cue (S1 in Table 1a), the first heuristic cue (H1:C1) and the first action of the first heuristic (H1:A1) are correct,³ that already ensures performing a correct action in 256 situations (25% of a total number of 1024 situations) and it is worth 25% of the overall fitness score (see Table 1b).

Given these dependencies, it is enough for a toolbox to have three actions and five cues correct in order to reach the 0.5 score of fitness (see Table 1b). The search space for mapping three actions to five cues is of size $50^3 \times 10^5 = 10^{10}$. This number holds given the assumption of equally distributed chances for a cue being true or false. In case one takes, say, a 1:10 ratio instead, the first action (H1:A1) is no longer worth 25% of fitness, but only 1%, which makes the search space grow drastically.

Part 2: Probabilities

Given the size of the search space for adaptive toolboxes, what is the probability that a random process—*à la* mutation and crossover—generates a toolbox of a certain level of fitness? To estimate these probabilities, we considered the fitness scores of any toolbox with cues and actions at each position of the toolbox being either correct or incorrect. Only a correct action can positively contribute to the overall fitness score of the toolbox. If all cues leading to this action are also correct, it increases the fitness by the relative probability of this action being executed. For example, if H1:A2 is correct and all of the cues S1, H1:C1 and H1:C2 are as well, the fitness of the toolbox is increased by the corresponding 12.5% points (see Table 1). However, if one of the cues leading to this action is incorrect, it will be executed in half of the cases. If two cues are incorrect, only in a quarter of the cases will the action be executed, and so on. Given the total number of actions and cues, the correct actions only occur in 2%, and correct cues in 10% of all possible toolboxes. That means that

³Note that, if for instance, the first heuristic cue (H1:C1, Table 1a) is incorrect (e.g., instead of C1, there is C3; and they are both either true or false), then it can still lead to execution of the first, and say, correct action (H1:A1). However, in half of the cases, where those cues are either true and false or false and true, that will not lead to execution of correct (H1:A1) action.

was 0.028, which is considerably lower than the 0.5 threshold that we defined for ecologically rational toolboxes. The fitness of the ‘best toolbox (from each generation) oscillates in the range [0.1, 0.4].

All simulations ended far before one thousand generations, often even before a hundred. All of the above indicate, that toolboxes perform poorly and do not improve with time. We explored two parameters which potentially could have influence the results. First, we reasoned that this effect might be due to a relatively low death rate. Such a low death rate (i) may ensure the survival and possibility of reproduction of individuals with lower fitness and (ii) imposes a lower pressure to select better toolboxes. Second, we explored the possibility of giving toolboxes more offspring. This change may lead to more populations surviving but we would not expect it to improve the overall individuals fitness. To test these predictions we ran two simulation studies, Conditions 2 and 3.

Condition 2: higher death rate

In this condition the death rate was increased ($p = 0.00045$; we opted for this relatively small increase in death rate, because a higher death rate would not afford successful runs, because none of individuals would survive the first survival-selection phase). In total, 20% of the simulations ended with a surviving population (Figure 4). The average performance of the surviving populations is 0.071, and the average performance for the dying out populations is 0.031. In order to calculate the average performance scores, we considered results from all the runs of simulations for surviving populations and all for the dying out populations separately (for a given condition), taking into account all the possible individual scores per every generation. Comparing the fitness in this Condition 2 with the fitness from Condition 1, it becomes clear that even if the higher pressure does improve performance of the toolboxes, as we had expected, the improvement is of a very small magnitude and does not bring the toolboxes anywhere closer to the 0.5 fitness.

Condition 3: higher chances of offspring

In this condition, the probability of generating a second child was increased to 47.4% per an individual. In total, 80% of the populations survived. As expected this survival rate was higher than in Condition 1 and 2. The average performance of the surviving populations is 0.044, and the average performance for the dying out populations is 0.027. In sum, the simulations in Condition 3 show that a larger growth rate leads to larger populations, but it does not make the individuals more ecologically rational.

Discussion

Using both formal argument and computer simulation, we have demonstrated the implausibility that phylogenetic processes (i.e., evolution) alone would ever produce ecologically rational adaptive toolboxes. Our simulations showed that populations of toolboxes that are ‘good enough’ to survive can evolve without these toolboxes showing any signs of

	Condition 1	Condition 2	Condition 3
% of Survival:	0%	20%	80%
Average Fit_S :	–	0.071	0.044
Average Fit_D :	0.028	0.031	0.027
Total average:	0.028	0.041	0.041

Table 3: Results from the simulations for the three different conditions (1: normal death rate; 2: higher death rate; 3: higher chances of offspring). Starting from the top, the rows show: percentage of surviving populations for every condition; the average fitness score (Fit_S) for a set of surviving populations per condition; the average fitness score (Fit_D) for a set of dying out populations per condition; the total average of a fitness score per condition.

‘ecological rationality’ (defined as the ability to make choices that are more often right than wrong; i.e. $\geq 50\%$ correct). In our simulation maximum fitness of populations hovered around 0.2 (20% correct decisions) and never got anywhere close to 0.5, let alone anything higher than that. The simulation results align well with our formal derivations: the expected number of generations needed to produce a toolbox grows exponentially. That means that even for only 10 possible cues and 50 possible actions the expected number of generations needed to produce at least one toolbox in the entire population with a fitness of at least 0.5 is 2,000,000 generations. For more possible cues or actions, the number of expected generations needed to produce at least one ecologically rational toolbox is even vastly larger.

Crucially, we refer here to the expected number of generations for producing a *single* toolbox with the feature of ‘ecological rationality’. Even if evolution would beat all odds and such an individual would be generated, the changes of its existence leading to a *population* with that feature are nanoscopically small. The reason is that toolboxes can survive with much lower fitness, and the chances of mutation and crossover leading to fitness below 0.2 is very high. With every new generation mutation and crossover occur, leading to a high probability that even if there is one ecologically rational individual in the pool that its offspring will be non-ecologically rational individuals that can again survive and procreate.

Does this mean that the adaptive toolbox account is implausible as an account of resource-bounded (human) decision making? Certainly not. Our findings do not rule out that adaptive toolboxes could be produced by ontogenic processes (learning and development), or even ontogenetic and phylogenetic processes combined (i.e., evolution could have produced those learning mechanisms that can produce adaptive toolboxes on a developmental time scale). After all, ontogenetic processes—unlike phylogenetic processes—are able to more actively search the space of possible parameters settings, e.g. by building a model of the environment and using that model to guide the search in a way that ensures ecologically rationality. However, in such a case it seems that one has




Figure 4: Examples of the simulations of dying out populations (a-c) and surviving populations (d-e). The plots show the changes of fitness over the time of many generations, including scores from the best (—), worst (—) and average (—) fitness per generation. For the normal death rate condition there is no surviving population.

to use a non-frugal learning mechanism to explain the emergences of adaptive toolboxes of fast and frugal heuristics. Resolving this tension seems an important target for future research in the area of resource-bounded decision making.

Acknowledgments

We would like to thank the Computational Cognitive Science group and in particular Mark Blokpoel for helpful insights and discussions. TW is supported by NSERC Discovery Grant RGPIN 228104-2010.

References

- Bergert, F. B., Nosofsky, R. M. (2007). A response-time approach to comparing generalized rational and take-the-best models of decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33(1).
- 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.
- Dieckmann, A., Rieskamp, J. (2007). The influence of information redundancy on probabilistic inferences. *Memory & Cognition*, 35(7):1801–1813.
- Gigerenzer, G. (2001). The adaptive toolbox: Toward a Darwinian rationality. In: D. W. Leger, A. C. Kamil, J. A. French (eds.), *Nebraska Symposium on Motivation*, vol. 47.
- Gigerenzer, G. (2002). The adaptive toolbox. In: G. Gigerenzer, R. Selten (eds.), *Bounded rationality: The adaptive toolbox*. MIT Press.
- Gigerenzer, G. (2004). Striking a blow for sanity in theories of rationality. *Models of a man: Essays in memory of Herbert A. Simon*, 389–409.
- Gigerenzer, G., Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62:451–482.
- Gigerenzer, G., Sturm, T. (2012). How (far) can rationality be naturalized? *Synthese*, 187(1):243–268.
- Gigerenzer, G., Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In: G. Gigerenzer, P. M. Todd, ABC Research Group (eds.), *Simple heuristics that make us smart*. Oxford University Press.
- Goldstein, D. G., Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In: G. Gigerenzer, P. M. Todd, ABC Research Group (eds.), *Simple heuristics that make us smart*. Oxford University Press.
- Martignon, L., Vitouch, O., Takezawa, M., Forster, M. R. (2003). Naive and yet enlightened: From natural frequencies to fast and frugal decision trees. In: D. Hardman, L. Macchi (eds.), *Thinking: Psychological perspective on reasoning, judgment, and decision making*. Wiley.
- Pohl, R. F. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making*, 19(3):251–271.
- Todd, P. M. (2002). Fast and frugal heuristics for environmentally bounded minds. In: G. Gigerenzer, R. Selten (eds.), *Bounded rationality: The adaptive toolbox*. The MIT Press Cambridge, MA.
- Todd, P. M., Gigerenzer, G. (1999). What we have learned (so far). In: G. Gigerenzer, P. M. Todd, ABC Research Group (eds.), *Simple heuristics that make us smart*. Oxford University Press.