

Cognitive Science 42 (2018) 1057–1065

Copyright © 2017 Cognitive Science Society, Inc. All rights reserved.

ISSN: 0364-0213 print / 1551-6709 online

DOI: 10.1111/cogs.12530

Demons of Ecological Rationality

Maria Otworowska,^{a,b} Mark Blokpoel,^{a,b} Marieke Sweers,^b
Todd Wareham,^c Iris van Rooij^{a,b}

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

^bDepartment of Artificial Intelligence, Radboud University

^cDepartment of Computer Science, Memorial University of Newfoundland

Received 21 March 2017; received in revised form 28 June 2017; accepted 14 July 2017

How can resource-bounded minds like our own make rational or otherwise “good” decisions in an uncertain and complex world (Oaksford & Chater, 1998; Simon, 1957, 1990)? The Adaptive Toolbox theory answers this question by defining human rationality in terms of a degree of adaptation of decision strategies (heuristics) to different environments (Gigerenzer & Todd, 1999; Todd & Gigerenzer, 2012). When heuristics are adapted to the environment and lead to “good enough” (or even high-quality) decisions, they are said to be *ecologically rational*. For almost two decades, this theory has been considered a tractable alternative to classical theories of human rationality based on logic or probability theory (Gigerenzer, 2015; Gigerenzer & Todd, 1999). These classical theories have been criticized for postulating intractable (e.g., NP-hard)¹ computations (Arkes, Gigerenzer, & Hertwig, 2016; Gigerenzer, 2008; Oaksford & Chater, 1998), which suggests that humans must possess demonic computational powers in order to make rational decisions (so-called *demons of rationality*; Gigerenzer & Todd, 1999; Goldstein & Gigerenzer, 1999). It is widely assumed that the Adaptive Toolbox theory circumvents the intractability problem that plagues classical accounts of human rationality, because heuristics are by definition tractable. Yet the notion of ecological rationality hinges on the existence of tractable adaptation processes. Here, we present an argument that, contrary to common belief, the Adaptive Toolbox theory has not yet tamed the intractability demon. Rather, the demon is hiding in the theory’s cornerstone assumption that ecological rationality is achieved by processes of adaptation, such as evolution, development, or learning.

Correspondence should be sent to Maria Otworowska, Donders Institute for Brain, Cognition, and Behaviour, Radboud University, Postbus 9104 6500 HE Nijmegen, The Netherlands. E-mail: m.otworowska@donders.ru.nl

The Adaptive Toolbox theory provides an influential account with many empirical successes (Brighton & Gigerenzer, 2012a,b; Bröder, 2000; Gigerenzer & Goldstein, 1996; Pohl, 2006; Schooler & Hertwig, 2005; Todd, 2001; Todd & Gigerenzer, 1999; 2012), which has led to its adoption in cognitive science, psychology, business, economics, law, philosophy, cultural studies, and medicine (Marewski & Gigerenzer, 2012; Todd & Gigerenzer, 2012). Despite its empirical successes, the theory remains incomplete to date (Todd & Gigerenzer, 2012). So far, research has focused on hypothesizing and testing the various heuristics in the toolbox, while two key aspects of the theory so far remained unresolved: (a) the meta-decision process of selecting the right heuristic for a given environment (the *selector*) (Hafenbrädl, Waeger, Marewski, & Gigerenzer, 2016; Todd & Brighton, 2016) and (b) the adaptation process by which the adaptive toolbox of heuristics evolves, develops, or is learned (Schulz, 2011).

First, several proposals about the nature of the selector have been suggested, but none so far is considered satisfactory (Marewski & Link, 2014). Be that as it may, it seems that to ensure tractability of the whole toolbox, minimally the selector must be fast and frugal like the heuristics that it selects (Gigerenzer & Todd, 1999, p. 32). Therefore—and to safeguard that our argument is not an artifact of a potentially intractable selector (cf. van Rooij, Wright, & Wareham, 2012)—we will work with the assumption that the selector itself is a heuristic as well. Second, the adaptation process involves creating and adapting the heuristics *and* the selector to be ecologically rational. It is assumed that both ontogenetic and phylogenetic adaptation processes can play a role (Todd & Brighton, 2016), but no explicit account of how this works has been put forth yet. To ensure generality of our result, we will make no assumption about the nature of the adaptation process other than that it yields toolboxes that are ecologically rational (cf. Otworowska et al., 2015).

Earlier work (Schmitt & Martignon, 2006) has already shown that *optimal* toolbox adaptation (defined as a problem of cue ordering in the toolbox) is intractable. These results were used as a supporting argument for the idea that ecological rationality is not defined in terms of optimality but in terms of “good enough” cue orders (Gigerenzer, 2008). This presupposes that “good enough” toolbox adaptation would be tractable. Here we show, however, that even “good enough” toolbox adaptation is intractable.

Importantly, intractability is not a property that can be derived from simulations, but given a proper formalization, it can be mathematically proven (van Rooij, 2008; van Rooij, Evans, Müller, Gedge, & Wareham, 2008). In the online supplementary materials,² we prove the intractability of toolbox adaptation. We first formalize the notions of a toolbox (heuristics + selector), ecological rationality, and the environment (Box 1). Then, using these notions, we formally define the Toolbox Adaptation problem (i.e., given an environment, create an ecologically rational [good enough] toolbox for that environment) (Box 1). Lastly, we construct a mathematical proof that Toolbox Adaptation, so defined, is intractable (NP-hard) (Box 2). Boxes 1 and 2 sketch properties of the formalization and give the intuition behind the proof.

Box 1: Formalization

To investigate claims about the computational (in)tractability of adaptive toolboxes, a formal, computational model of toolbox adaptation is required. Such a model can be represented by an input-output mapping (i.e., what Marr (1982) would refer to as “computational level model”; see also van Rooij, 2008) as follows:

Input: An environment, that is, a set of actions, and a set of situations (formalized as truth assignments for possible events). An upperbound on the number of heuristics ($\#h$) and the size of a heuristic ($|h|$). A lowerbound for the level of adaptation that counts as ecologically rational ($er_{min} \in [0, 1]$).

Output: A toolbox T , of bounded size, that is ecologically rational.

Fig. 1 illustrates a possible (toy example) input and output for Toolbox Adaptation. Note that in this computational-level model, the toolbox T consists only of fast and frugal trees (i.e., both the selector and heuristics are fast and frugal trees). A fast and frugal tree is a chain of cues with associated actions (in case of a heuristic) or a chain of cues with associated heuristics (in case of the selector). Each cue $c(e, s)$ is a boolean function, evaluating whether an event ($e \in E$) is true in a given situation. If the cue $c(e, s)$ evaluates to true, then the heuristic associated with that selector cue is used (in case of the selector) or an action associated with that heuristic cue is executed (in case of a heuristic). If the cue is false, the next cue is evaluated until the last cue is reached. If this last cue is false, the first heuristic is used (in case of the selector) or the last action in the tree is performed (in case of a heuristic).

The choice for fast and frugal trees is without loss of generality, because (a) many other heuristics proposed for the Adaptive Toolbox theory, such as for example, fluency heuristic, take the Best, satisficing, $1/N$, default heuristic, tit-for-tat, imitate the majority, and imitate the successful (Gigerenzer, 2008), can be formally rewritten as fast and frugal trees (Sweers, 2015); and (b) if adaptation of toolboxes is intractable for some subset of heuristics, then it is also intractable for any superset of that.

Given that our computational model of Toolbox Adaptation is an input-output mapping, it is neutral with respect to the nature of the adaptation process by which the output is reached. For instance, this process could be an ontogenetic or phylogenetic process, or a mixture of these. Furthermore, it does not make specific assumptions about how these processes are realized, for example, algorithmically. The results of (in)tractability analyses of a model like this will therefore hold for any type of algorithmic-level implementation, which could be either evolutionary, neural network, probabilistic, incremental, hill climbing, or any other type of algorithm. The reason is that computational intractability (i.e., NP-hardness) is a property of the input-output mapping, and not of a specific algorithm for computing it (Garey & Johnson, 1979).




Fig. 1. Formalization of toolbox adaptation.

The NP-hardness proof for Toolbox Adaptation establishes that there does not exist any general polynomial-time computable process (neither deterministic nor probabilistic³; see also van Rooij, 2008) that can adapt toolboxes to be ecologically rational (“good enough”), for all possible environments. This applies regardless of the nature of this process.⁴ More important, it demonstrates that Toolbox Adaptation is as difficult to compute as many other known NP-hard functions, including logic problems, such as deciding logical satisfiability of a set of logical clauses (Gary & Johnson, 1979; Oaksford & Chater, 1998), and probabilistic inference problems, such as exact or approximate inference in Bayesian networks (Abdelbar, Hedetniemi, & Hedetniemi, 2000; Kwisthout, Wareham, & van Rooij, 2011). This is an interesting observation given that one of the prime motivations for the Adaptive Toolbox theory was to move away from classical notions of rationality, based on logic or probability, in order to ensure tractability.

Our proof that Toolbox Adaptation is intractable may be surprising, given that it is so widely believed that the Adaptive Toolbox theory is a tractable account of human rationality. We suspect that the belief could persist, however, because researchers have been focusing on Toolbox *Application*, while taking Toolbox *Adaptation* for granted. Here, Toolbox Application refers to the process of making ecologically rational decisions in a given environment, using a toolbox of heuristics that has already been adapted to that environment by some unspecified process. Even if Toolbox Application is free from computational demons, the demons are still hiding in Toolbox Adaptation.

It is not uncommon for cognitive scientists to try to discredit theories in competing frameworks by pointing out that those frameworks run into intractability issues. But this is to no avail and is in no way our purpose here. We see intractability not as a problem for specific theories, or even for specific theoretical frameworks, but a ubiquitous feature of theoretical frameworks with high degrees of generality (van Rooij, 2008, 2015). For instance, Bayesians originally criticized logical accounts of rationality for their intractability (Chater & Oaksford, 1993; Oaksford & Chater, 1998), only to later discover that Bayesian theories themselves face intractability charges that are not easily fenced off by appeals to “approximation” or “as if” explanation (Kwisthout et al., 2011; van Rooij, Wright, Kwisthout, & Wareham, 2014). Similarly, Gigerenzer and colleagues have criticized both logical and Bayesian accounts of rationality for their intractability (Gigerenzer & Todd, 1999; Todd, 2001). By overlooking the question how adaptation of toolboxes of heuristics can itself be tractable, Gigerenzer and colleagues may not have realized that the Adaptive Toolbox theory faces exactly the same intractability charge, albeit in a different guise. From a complexity-theoretic perspective this is not surprising, but a natural consequence of the theory’s high degrees of expressiveness (i.e., it has the degrees of freedom needed to encode NP-hard problems).

Does this mean that the Adaptive Toolbox account is irreparably flawed? Certainly not. In fact, there exists a well-defined methodology with which it could be repaired. Using proof techniques from parameterized complexity theory (Downey & Fellows, 1999), it is possible to identify parameters of an intractable model that are sources of the intractability, in the sense that they can confine the non-polynomial time complexity inherent to a model (van Rooij & Wareham, 2008; van Rooij et al., 2008). The model

Box 2: Proof intuition

NP-hard models are considered to be intractable, because there do not exist any tractable (formally, polynomial-time) algorithms for computing them (unless $P = NP$, see footnote 1). How can one prove that a computational model is NP-hard (and thus intractable)? A model is NP-hard if it is at least as hard as another NP-hard model, that is, one can use the former to compute the latter with minimal (only polynomial) overhead. We can prove such a relationship using polynomial-time reduction, a technique from computational complexity (Garey & Johnson, 1979). In the supplementary materials, we use this technique to prove that Adaptive Toolbox is NP-hard. Specifically, we reduce from the known NP-hard graph problem Dominating Set (which takes as input any graph and asks as output a so-called dominating set, that is, a subset of vertices with the property that for each vertex v in the graph, either v or one of its neighbors is in the subset). This is done as follows:

- Provide a tractable (polynomial-time) algorithm that transforms any input i_{DS} for Dominating Set into input i_{AT} for Toolbox Adaptation (proof step 1 in Supplementary Materials).
- Show that after the transformed input is processed as specified by Toolbox Adaptation, the resulting output o_{AT} can be transformed back into a solution o_{DS} for Dominating Set (proof steps 2 and 3).
- Both transformations only take polynomial time to compute. This means we can use Toolbox Adaptation to solve Dominating Set with minimal (only polynomial) overhead.
- Toolbox Adaptation is NP-hard, because it is at least as hard as the NP-hard problem Dominating Set.

This is, in a nutshell, the strategy we used to prove that Toolbox Adaptation is NP-hard. Fig. 2 illustrates this strategy.

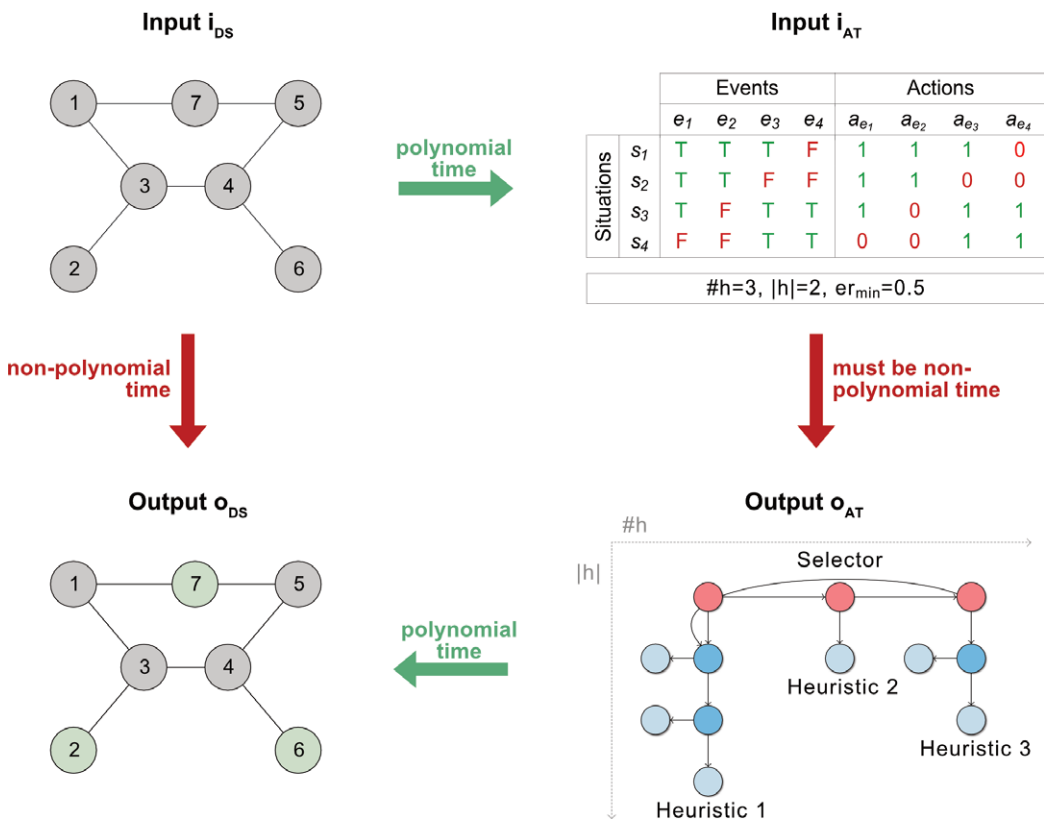


Fig. 2. Proof intuition.

can then potentially be rendered tractable by assuming that the relevant parameters are, in the real world, bounded in the right ways. In other words, these bounds limit the degrees of freedom of the theory so that it becomes tractable for ecologically relevant parameter ranges. Importantly, taking this approach would require that Adaptive Toolbox theorists extend their research program with a “tractable design cycle” (analogous to, and interlaced with, the “empirical cycle”; cf. van Rooij, 2008, p. 962):⁵

1. Consider parameters of the model (these can be one or more properties of the toolbox, the environment, or ecological rationality).
2. Use formal proof techniques to analyze if these parameters are sources of intractability for the model.
3. **If Yes, then** consider if the parameters can plausibly be assumed to be bounded in real-world situations (not just in toy domains).
4. **If Yes, then** revise the model by explicating bounds on these parameters.
5. **Else**, return to step 1.
6. **Else**, return to step 1.

Adopting this methodology not only has the benefit that it can potentially render a tractable version of the Adaptive Toolbox theory, but it may also sharpen the debate among logicians, Bayesians, and heuristics. After all, classical approaches to rationality have the same methodology at their disposal (see also Kwisthout et al., 2011; van Rooij et al., 2014; van Rooij and Wareham, 2008). Applying the “tractable design cycle” to both ecological and classical accounts of rationality is a rigorous way to move forward on the question how rationality can be “tractable in the real world in which people live, not only in the small world of an experiment” (Gigerenzer et al., 2008, p. 236), as well as to assess whether or not the ecological account can really explain this better than classical accounts.

Acknowledgments

The authors thank Richard Cooper and two anonymous reviewers for their useful comments. M.O. was supported by a Donders Centre for Cognition PhD grant awarded to I.v.R., and T.W. was supported by National Science and Engineering Research Council (NSERC) Discovery grants 228104-2010 and 228104-2015.

Notes

1. Given the famous mathematical conjecture that $P \neq NP$ (Fortnow, 2009), NP-hard problems cannot be computed in polynomial time, but instead require non-polynomial (e.g., exponential) time (Arora & Barak, 2009; Gary & Johnson, 1979). Here, P is the class of problems that are efficiently (polynomial time) solvable, and NP is the class for which solutions are efficiently verifiable. Because of the combinatorial explosion inherent in non-polynomial computations, NP-hard computations are

generally regarded intractable and therefore models that postulate such functions are considered computationally implausible (van Rooij, 2008). Or as Gigerenzer and colleagues have put it: “The computations postulated by a model of cognition need to be tractable in the real world in which people live, not only in the small world of an experiment with only a few cues. This eliminates NP-hard models that lead to computational explosion, such as probabilistic inference using Bayesian belief networks (Cooper, 1990), including its approximations (Dagum & Luby, 1993). Tractability is one condition for psychological plausibility that is easily met by simple heuristics” (Gigerenzer, Hoffrage, & Goldstein, 2008, p. 236).

2. <http://www.dcc.ru.nl/irisvr/papers/suppl17.pdf>.
3. The notion of probabilistic algorithms is captured by the class BPP of problems which have bounded-error probabilistic polynomial-time algorithms. It is assumed that the class BPP is equal to P, so that functions which can be solved probabilistically in polynomial time can also be solved deterministically in polynomial time (Wigderson, 2006). Hence, NP-hard functions cannot be solved (with bounded error) by any probabilistic polynomial-time algorithm (assuming $P \neq NP$; see also note 1).
4. Given that NP-hardness is a property of the function computed and not of the algorithm used to compute the function, the NP-hardness result applies to any such process. See also Box 2.
5. This cycle could be implemented in practice using concepts and techniques from parameterized complexity theory. It can be shown that some intractable (NP-hard) functions $f : I \rightarrow O$ can be computed in fixed-parameter (fp-) tractable time $O(g(K)lil^c)$, that is, where g can be any function of the parameters k_1, k_2, \dots, k_m in set $K = \{k_1, k_2, \dots, k_m\}$, lil denotes the input size, and c is a constant. Note that in such an event, the intractable f can be computed efficiently (in polynomial time), even for large inputs, provided the assumption that f operates only on inputs in which the parameters in K are restricted to relatively small values (each $k \ll lil$). The use of such formal proof techniques is no luxury, as it is known that intuitions about what are (and are not) the sources of intractability in a model can often be mistaken (van Rooij et al., 2008). For instance, one may intuitively conjecture that Toolbox Adaptation is tractable for toolboxes of bounded size (e.g., a small [fixed] number of heuristics available in the toolbox), but a corollary of our proof is that this is false. In the reduction in our proof the parameter $\#h$ is set to 1, which shows that Toolbox Adaptation is NP-hard even if the toolbox would contain only a single heuristic.

References

- Abdelbar, A. M., Hedetniemi, S. T., & Hedetniemi, S. M. (2000). The complexity of approximating maps for belief networks with bounded probabilities. *Artificial Intelligence*, 124(2), 283–288.
- Arkes, H. R., Gigerenzer, G., & Hertwig, R. (2016). How bad is incoherence? *Decision*, 3(1), 20–39.

- Arora, S., & Barak, B. (2009). *Computational complexity: A modern approach*. Cambridge: Cambridge University Press.
- Brighton, H., & Gigerenzer, G. (2012a). Homo heuristicus: less-is-more effects in adaptive cognition. *The Malaysian journal of Medical Sciences: MJMS*, 19(4), 6.
- Brighton, H. J., & Gigerenzer, G. (2012b). How heuristics handle uncertainty. In P. M. Todd & G. Gigerenzer (Eds.), *Ecological rationality: Intelligence in the world* (pp. 33–60). Oxford, UK: Oxford University Press.
- 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.
- Chater, N., & Oaksford, M. (1993). Logicism, mental models and everyday reasoning: Reply to Garnham. *Mind & Language*, 8(1), 72–89.
- Cooper, G. F. (1990). The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2–3), 393–405.
- Dagum, P., & Luby, M. (1993). Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artificial Intelligence*, 60(1), 141–153.
- Downey, R. G., & Fellows, M. R. (1999). *Parameterized complexity*. New York: Springer.
- Fortnow, L. (2009). The status of the P versus NP problem. *Communications of the ACM*, 52(9), 78–86.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness* (Series of Books in the Mathematical Sciences). W. H. Freeman & Co, New York, p. 340. <https://doi.org/10.1137/1024022>
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. New York: WH Freeman and Company.
- Gigerenzer, G. (2008). Why heuristics work. *Perspectives on Psychological Science*, 3(1), 20–29.
- Gigerenzer, G. (2015). *Simply rational: Decision making in the real world*. Oxford, UK: Oxford University Press.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650.
- Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In G. Gigerenzer, P. M. Todd, and the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 3–34). Oxford, UK: Oxford University Press.
- Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins, and Thomas (2008). *Psychological Review*, 115(1), 230–239.
- Goldstein, D. G., & Gigerenzer, G. (1999). The recognition heuristic: How ignorance makes us smart. In G. Gigerenzer, P. M. Todd, and the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 37–58). Oxford, UK: Oxford University Press.
- Hafenbrädl, S., Waeger, D., Marewski, J. N., & Gigerenzer, G. (2016). Applied decision making with fast-and-frugal heuristics. *Journal of Applied Research in Memory and Cognition*, 5(2), 215–231.
- Kwisthout, J., Wareham, T., & van Rooij, I. (2011). Bayesian intractability is not an ailment that approximation can cure. *Cognitive Science*, 35(5), 779–784.
- Marewski, J. N., & Gigerenzer, G. (2012). Heuristic decision making in medicine. *Dialogues in Clinical Neuroscience*, 14(1), 77–89.
- Marewski, J. N., & Link, D. (2014). Strategy selection: An introduction to the modeling challenge. *Wiley Interdisciplinary Reviews: Cognitive Science*, 5(1), 39–59.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: Henry Holt and Co.
- Oaksford, M., & Chater, N. (1998). *Rationality in an uncertain world: Essays on the cognitive science of human reasoning*. Abingdon, UK: Psychology Press/Erlbaum (UK) Taylor & Francis.
- Otworowska, M., Sweers, M., Wellner, R., Uhlmann, M., Wareham, T., & van Rooij, I. (2015). How did Homo Heuristicus become ecologically rational? In G. Airenti, B. G. Bara, & G. Sandini, (Eds.), *EuroAsianPac Joint Conference on Cognitive Science* (pp. 324–329). Torino, Italy. ceur-ws.org/Vol-1419/

- Pohl, R. F. (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making*, 19(3), 251–271.
- Schmitt, M., & Martignon, L. (2006). On the complexity of learning lexicographic strategies. *Journal of Machine Learning Research*, 7(Jan), 55–83.
- Schooler, L. J., & Hertwig, R. (2005). How forgetting aids heuristic inference. *Psychological Review*, 112(3), 610.
- Schulz, A. W. (2011). Gigerenzer's evolutionary arguments against rational choice theory: An assessment. *Philosophy of Science*, 78(5), 1272–1282.
- Simon, H. A. (1957). *Models of man: Social and rational*. New York: John Wiley and Sons, Inc.
- Simon, H. A. (1990). Bounded rationality. In J. Eatwell, M. Milgate, P. Newman, (Eds.), *Utility and probability* (pp. 15–18). London: The New Palgrave.
- Sweers, M. (2015). Adapting the adaptive toolbox: The computational cost of building rational behavior. Master's thesis, Radboud University Nijmegen, the Netherlands.
- Todd, P. M. (2001). Fast and frugal heuristics for environmentally bounded minds. In R. Selten, & G. Gigerenzer, (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 51–70). Cambridge, MA: MIT Press.
- Todd, P. M., & Brighton, H. (2016). Building the theory of ecological rationality. *Minds and Machines*, 26(1–2), 9–30.
- Todd, P. M., & Gigerenzer, G. (1999). What we have learned (so far). In G. Gigerenzer, P. M. Todd, and the ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 357–365). Oxford, UK: Oxford University Press.
- Todd, P. M., & Gigerenzer, G. (2012). What is ecological rationality? In P. M. Todd, & G. Gigerenzer, (Eds.), *Ecological rationality: Intelligence in the world* (pp. 3–30). Oxford, UK: Oxford University Press.
- van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, 32(6), 939–984.
- van Rooij, I. (2015). How the curse of intractability can be cognitive science's blessing. In R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, P. P. Maglio, & D. C. Noelle (Eds.), *Proceedings of the 37th Annual Meeting of the Cognitive Science Society* (pp. 2839–2840). Austin, TX: Cognitive Science Society.
- van Rooij, I., & Wareham, T. (2008). Parameterized complexity in cognitive modeling: Foundations, applications and opportunities. *The Computer Journal*, 51(3), 385–404.
- van Rooij, I., Evans, P., Müller, M., Gedge, J., & Wareham, T. (2008). Identifying sources of intractability in cognitive models: An illustration using analogical structure mapping. In B. C. Love, K. McRae, and V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 915–920). Austin, TX: Cognitive Science Society.
- van Rooij, I., Wright, C. D., & Wareham, T. (2012). Intractability and the use of heuristics in psychological explanations. *Synthese*, 187(2), 471–487.
- van Rooij, I., Wright, C., Kwisthout, J., & Wareham, T. (2014). Rational analysis, intractability, and the prospects of “as if”-explanations. *Synthese*. <https://doi.org/10.1007/s11229-014-0532-0>.
- Wigderson, A. (2006). P, NP and mathematics—A computational complexity perspective. In M. Sanz-Solé, J. Soria, J. Luis Varona, J. Verdera, (Eds.), *Proceedings of the 2006 International Congress of Mathematicians* (pp. 665–712). Madrid: European Mathematical Society Publishing House.

Supporting Information

Additional Supporting Information may be found online in the supporting information tab for this article:

Appendix S1: Supplementary Materials.