

Toward a Formal Science of Heuristics

Ardavan S. Nobandegani^{1,3} & Thomas R. Shultz^{2,3}

{ardavan.salehinobandegani, william.campoli}@mail.mcgill.ca
{thomas.shultz}@mcgill.ca

¹Department of Electrical & Computer Engineering, McGill University

²School of Computer Science, McGill University

³Department of Psychology, McGill University

Abstract

Heuristics are simple, effective cognitive processes that deliberately ignore parts of information relevant to decision-making. Ecological rationality, as an essential part of the Adaptive Toolbox research program on heuristics, investigates the environmental conditions under which simple heuristics would outperform complex models of decision-making, thereby providing support for the surprising less-is-more effect. In this work, we present a new research program, dubbed formal science of heuristics (FSH), that nicely complements the ecological rationality research, developing it into a much richer research program. Concretely, FSH sets to (i) mathematically delineate the broadest class of environmental conditions under which a heuristic is fully optimal, and (ii) formally investigate how deviations from those conditions would lead to degradation of performance, thereby allowing for a mathematically rigorous characterization of their robustness. As an instantiation of the FSH research program, we present several analytical results aiming to delineate the mildest conditions granting the optimality of a well-known heuristic: Take The Best. We conclude by discussing the implications that pursuit of FSH could have on the science of heuristics.

Keywords: Ecological rationality; one-reason heuristics; formal science of heuristics; Take The Best heuristic

1 Introduction

Heuristics—simple, effective cognitive processes that deliberately ignore parts of information relevant to decision-making—are assumed to underpin much of human judgment and decision-making (e.g., Gigerenzer & Selten, 2001; Mousavi, Gigerenzer, & Kheirandish, 2016), and are widely considered to be sub-optimal, attaining higher speed at the expense of lower accuracy (e.g., Payne et al., 1993; Shah & Oppenheimer, 2008; Evans and Over, 2010).

Challenging the latter mindset, the influential *ecological rationality* research program (as part of the Adaptive Toolbox theory) maintains that heuristic are well-matched to the environment they are adopted in (Todd & Gigerenzer, 2007), and seeks to investigate the environmental conditions under which heuristics would outperform complex models of decision-making, giving rise to the surprising less-is-more effect: when less information or computation leads to more accurate judgments than more information or computation (Gigerenzer & Gaissmaier, 2011).

Despite great successes, ecological rationality work has predominantly focused on simulation-based demonstrations of simple heuristics outperforming complex strategies (e.g., Gigerenzer et al., 2008, Gigerenzer & Todd, 1999, Todd & Gigerenzer, 2000, Hoffrage & Reimer, 2004; Gigerenzer & Goldstein, 1996), directing comparatively little effort (but

see, e.g., Martignon & Hoffrage, 2002, Hogarth & Karelaia, 2006) toward establishing a mathematically-rigorous characterization of the environmental conditions underpinning the less-is-more effect — Todd and Gigerenzer (2007) explicitly call for developing such deep theoretical accounts.

In this work, we present a new research program, dubbed formal science of heuristics (FSH), that nicely complements the ecological rationality research, developing it into a richer research program, and, additionally, permitting mathematicians and computer scientists to make important contributions to the science of heuristics.

Concretely, FSH pursues the following two objectives. (1) FSH seeks to mathematically delineate the *broadest* class of environmental conditions under which a heuristic is fully optimal (i.e., using the standard terminology of computer science, the environmental conditions under which a heuristic serves as a correct algorithm w.r.t. the objective of interest, or, equivalently, an approximation algorithm with an approximation ratio of one). (2) FSH aims to formally investigate how deviations from optimality conditions would lead to degradation of performance, thereby allowing for a mathematically rigorous characterization of a heuristic’s robustness. As such, to provide strongest theoretical support for the robustness of a heuristic, FSH aims to analytically provide the *mildest* technical conditions granting the optimality of a heuristic. According to the Adaptive Toolbox theory, robustness is predominantly responsible for the less-is-more effect, and plays a central role in the success of fast-and-frugal heuristics in everyday life decisions (e.g., Gigerenzer & Todd, 1999; Gigerenzer & Gaissmaier, 2011).

With regard to objective (2) mentioned above, one of the mildest technical conditions worth considering is *distribution-free* performance guarantees, widely studied in statistical learning theory and machine learning (e.g., Valiant, 1984; Kearns, Vazirani, & Vazirani, 1994). As is often the case, a decision-maker lacks (at least partially) the knowledge of the regularities of their environment, and, therefore, is not fully informed as to how information relevant to a decision-making task of interest is distributed. Distribution-free results, as the term suggests, establish performance guarantees that hold true *regardless* of the probability distribution governing a decision-making task (e.g., the distribution of attributes in a multi-alternative decision-making task). As such, distribution-free results provide strong robustness guarantees while demanding *minimal* environmental knowledge on the

part of the decision-maker, thus playing an integral role in the FSH research program.

We should note that establishing distribution-free performance guarantees for a heuristic does *not* imply that: (1) the decision-maker is inattentive to their environment, nor that (2) the decision-maker is not trying to select a heuristic well-matched to the environment — experimental evidence clearly suggests otherwise (e.g., Rieskamp & Otto, 2006; Hoffart, Rieskamp, & Dutilh, 2018; Payne, Bettman, & Johnson, 1988; Bröder, 2003; Pachur, Todd, Gigerenzer, Schooler, & Goldstein, 2011). On the contrary, establishing distribution-free performance guarantees on a heuristic ensures that that heuristic is well-matched to the environment, even when the decision-maker’s knowledge of the environment is *imperfect*—a psychologically plausible assumption.

This work is organized as follows. We begin by presenting an overview of a well-known heuristic: Take The Best (TTB). As an instantiation of FSH, we then establish several analytical results, including strong distribution-free performance guarantees, for TTB. Finally, we conclude by discussing the implications that pursuit of FSH could have on the science of heuristics.

2 Take The Best: An Overview

Take The Best (TTB; Tversky, 1969, Gigerenzer, Hoffrage, & Kleinbölting, 1991) belongs to the class of one-reason decision-making heuristics which base decisions on only one attribute value. In its classic form, TTB is concerned with the task of predicting which of two objects, each possessing several binary-valued attributes, has a higher value on a given criterion, e.g., which of two cities has a higher population, or, which of two cookies would be more delicious.

The machinery of TTB is quite simple: Starting with the attribute having the highest validity, make pairwise-comparisons between the attribute values of the two objects; as soon as the first discriminating attribute is encountered (i.e., the attribute on which the two objects differ), announce the object attaining the highest attribute value on the discriminating attribute to be the winning object. TTB visits attributes in a descending order of their validities. If no discriminating attribute is ever encountered, TTB selects the winning object uniformly at random.¹ In TTB, the validity v_i of the i^{th} attribute is given by (Gigerenzer et al., 2008):

$$v_i := \frac{R_i}{R_i + W_i},$$

where R_i, W_i are the number of correct and incorrect inferences based on the i^{th} attribute alone, respectively.

The efficacy of TTB receives strong empirical support from a wide range of economic, demographic, environmental, and other prediction tasks (e.g., Gigerenzer et al., 2008; Czerlinski, Gigerenzer, & Goldstein, 1999; Chater, Oaksford,

¹Without loss of generality, we assume that the decision-maker initially recognizes all the objects which s/he has to choose from. Accordingly, the use of recognition heuristic (Gigerenzer et al., 2008), as the first step of TTB, is implicitly considered in our work.

Nakisa, & Redington, 2003). For example, on the task of predicting which of two cities has a higher homeless rate, TTB achieves better prediction accuracy than several competitors, including multiple regression model (Gigerenzer et al., 2008). More strikingly, Czerlinski, Gigerenzer, and Goldstein (1999) empirically showed that, across 20 real-world prediction problems, on average TTB obtains the best prediction accuracy when competing with several prominent alternatives, including multiple regression and tallying heuristic. Relatedly, on the same 20 real-world prediction problems, Gigerenzer et al. (2008) empirically show that the predictive accuracy of TTB came, on average, within three percentage points of a complex Bayesian network model. More broadly, when environments are moderately unpredictable and learning samples are small, as with many social and economic situations, TTB tends to make inferences as accurately as or better than multiple regression and neural networks (Chater, Oaksford, Nakisa, & Redington, 2003).

To provide direct experimental evidence for TTB as a psychological model, Bröder and his colleagues (Bröder 2000; Bröder and Schiffer 2003) conducted 20 studies, concluding that TTB is used under a number of conditions such as when information is costly and the variability of the validity of the attributes is high. Furthermore, Bröder and Gaissmaier (2007) and Nosofsky and Bergert (2007) showed that TTB predicts response times better than weighted additive and exemplar models.

Previous work assessing the prediction accuracy of TTB has mainly focused on computer simulations, with some work establishing analytical results formally supporting the efficacy of TTB (e.g., Martignon & Hoffrage, 2002; Hogarth & Karelaia, 2005, 2006; Baucells, Carrasco, & Hogarth, 2008).

Pursuing the research program proposed by FSH, and contrary to past analytical work, in this work we consider a much broader class of problems involving nonlinear objective functions (Definitions 1-4) with interactions between attributes being also accounted for (Definition 2). For the broad class of problems discussed above, we formally establish conditions granting the optimality of TTB when dealing with both non-binary, discrete attribute values (Propositions 3, 5, and 6) and continuous attribute values (Proposition 4). We also analytically investigate a broad class of prediction problems—involving both structured (Definition 3) and unstructured noise (Definition 4)—for which only *probabilistic* guarantees can be provided. Additionally, and in sharp contrast to past analytical work, we provide strong distribution-free guarantees on TTB for several classes of prediction problems (Propositions 5, 6, and 8).

3 Instantiating FSH: The Case of TTB

As an instantiation of FSH, in this section we establish several analytical results, including strong distribution-free performance guarantees, for TTB.

Before we proceed further, let us formally delineate an objective function which characterizes a broad class of decision-

making problems.

Definition 1. (Objective function) Let O_1, O_2, \dots, O_N denote the set of N objects a decision-maker should choose from. Let also $O_{ij} \in \{0, 1\}$ denote the value of the j^{th} attribute of object O_i where $1 \leq j \leq M$, and w_k denote the weight corresponding to the k^{th} attribute, A_k . Finally, let $\psi(\cdot)$ be an arbitrary monotonically-increasing function (i.e., $\forall x: \frac{d}{dx}\psi(x) > 0$). Then, the winning object O_{i^*} is the one whose index i^* satisfies the following objective function:

$$i^* := \arg \max_i \psi\left(\sum_{j=1}^M w_j O_{ij}\right). \quad (1)$$

Therefore, in the case of having only two objects O_1, O_2 to choose from, the optimal decision rule is given by:

$$\psi\left(\sum_{j=1}^M w_j O_{1j}\right) \underset{O_1}{\overset{O_2}{\succ}} \psi\left(\sum_{j=1}^M w_j O_{2j}\right), \quad (2)$$

where $A \underset{O_1}{\overset{O_2}{\succ}} B$ denotes the following: choose O_1 if $A > B$; choose O_2 if $A < B$; and choose uniformly at random between O_1, O_2 if $A = B$. ■

It is worth noting that in Eqs. (1-2), ψ can be any monotonically-increasing function, e.g., $\psi(x) = e^x, \psi(x) = 2^x + \log(x)$.

Proposition 1. (Sufficient condition for optimality) If there exists a $k \in \mathbb{N}$ such that $\forall p < k, \forall i, j \in \{1, \dots, N\} O_{ip} = O_{jp}$ and $w_k > \sum_{i>k} w_i$, then the following holds true:

$$\exists j \in \{1, \dots, N\} \forall i \neq j O_{jk} > O_{ik} \Rightarrow O_{i^*} := O_j. \quad (3)$$

Importantly, Proposition 1 establishes a condition granting basing decision on only one attribute while preserving optimality with respect to the objective given in (1). As such, Proposition 1 provides a firm rational basis for the possibility of one-reason decision-making for the broad class of decision-making problems characterized in Definition 1.

Next, Proposition 2 establishes a condition granting the optimality of TTB (when choosing between an arbitrary number of objects) with respect to the objective given in (1).

Proposition 2. (Generalizing TTB to N -object prediction tasks) Let O_1, O_2, \dots, O_N denote the set of N objects a decision-maker is to choose from. Let also w_k denote the weight corresponding to the k^{th} attribute (see Definition 1), and v_k denote the validity of the k^{th} attribute. If $\forall k w_k = v_k$ and $\exists r \in \mathbb{R}^{>2}$ s.t. $\forall k v_k \leq \left(\frac{1}{r}\right)v_{k-1}$, then TTB is an optimal strategy for the class of decision-making problems characterized in Definition 1. ■

In the N -object setting (as in Proposition 2), TTB works as follows: Starting with the attribute having the highest validity, compare attribute values across the N objects; as soon as the first discriminating attribute is encountered (i.e., the attribute on which at least two objects differ), exclude from consideration those objects faring worse on the discriminating attribute; announce the object surviving this elimination

process to be the winning object. TTB visits attributes in a descending order of their validities. If no discriminating attribute is ever encountered, TTB selects the winning object uniformly at random.

Proposition 3. (Multi-level attribute values) Let O_1, O_2, \dots, O_N denote the set of N objects a decision-maker is to choose from, with each object having M attributes. Let also $O_{ij} \in \{0, 1, \dots, \theta\}$ denote the value of the j^{th} attribute of object O_i where $1 \leq j \leq M$. Finally, let w_k denote the weight corresponding to the k^{th} attribute (see Definition 1), and v_k denote the validity of the k^{th} attribute. If $\forall k w_k = v_k$, and $\exists r > 1 + \theta$ s.t. $\forall k v_k \leq \left(\frac{1}{r}\right)v_{k-1}$, then TTB is an optimal strategy for the class of decision-making problems characterized in Definition 1. ■

In simple terms, Proposition 3 analytically establishes a conditions granting the optimality of TTB (when generalized to the setting of N objects, each with discrete, multi-level attribute values) with respect to the objective given in (1).

Proposition 4. (Continuous attribute values) Let $\forall i \in \{0, 1\}, O_i$ denote the two objects a decision-maker should choose from, with each object having M attributes. Let also $O_{ij} \in \mathbb{R}$ denote the value of the j^{th} attribute of object O_i . Finally, let w_k denote the weight corresponding to the k^{th} attribute (see Definition 1), and v_k denote the validity of the k^{th} attribute. Assuming that k^* denotes the index of the discriminating attribute on which TTB halts, and $|O_{1k^*} - O_{2k^*}| \leq \delta$, the following statement holds true: If $\forall k w_k = v_k$, and $\forall i, j \in \{1, \dots, M\} O_{ij} \leq U$, and $\exists r > 1 + \frac{U}{\delta}$ s.t. $\forall k v_k \leq \left(\frac{1}{r}\right)v_{k-1}$, then TTB is an optimal strategy for the class of decision-making problems characterized in Definition 1.

Proposition 4 analytically establishes a conditions granting the optimality of TTB (when choosing between two objects, each with continuous attribute values) with respect to the objective given in (1).

Following the line of research proposed by FSH, next we present our first distribution-free guarantee for TTB.

Proposition 5. (Distribution-free guarantee) Let O_1, O_2, \dots, O_N denote the set of N objects a decision-maker is to choose from, with each object having M attributes. Let also $O_{ij} \in \{0, 1, \dots, \theta\}$ denote the value of the j^{th} attribute of object O_i , where $\{O_{ij}\}_{i,j} \stackrel{d}{\sim} P$ with P denoting a joint probability distributions over the set of all attribute values $\{O_{ij}\}_{i,j}$. Finally, let w_k denote the weight corresponding to the k^{th} attribute (see Definition 1), and v_k denote the validity of the k^{th} attribute. Then, for any joint probability distribution P the following statement holds true: If $\forall k w_k = v_k$, and $\exists r > 1 + \theta$ s.t. $\forall k v_k \leq \left(\frac{1}{r}\right)v_{k-1}$, then TTB is an optimal strategy for the class of decision-making problems characterized in Definition 1. ■

Proposition 5 analytically establishes a condition ensuring the optimality of TTB (when generalized to the setting of N objects, each with discrete, multi-level attribute values) with respect to the objective given in (1), in a strong distribution-

free manner. It is crucial to note that the optimality guarantee given in Proposition 5 holds true for *any* joint distribution P on the set of all attribute values.

Next, in Definition 2, we formally characterize a broad class of prediction problems wherein interactions between attribute values are also accounted for.

Definition 2. (Objective function) Let O_1, O_2, \dots, O_N denote the set of N objects a decision-maker should choose from, and $O_{ij} \in \{0, 1, \dots, \theta\}$ denote the value of the j^{th} attribute of object O_i where $1 \leq j \leq M$. Additionally, let w_k denote the weight corresponding to the k^{th} attribute, and r_{pq} denote the weight quantifying the amount of interaction between the p^{th} and the q^{th} attributes. Finally, let $\chi(\cdot)$ be an *arbitrary* monotonically-increasing function (i.e., $\forall x: \frac{d}{dx}\chi(x) > 0$). Then, the winning object O_{i^*} is the one whose index i^* satisfies the following objective function:

$$i^* \triangleq \arg \max_i \chi \left(\sum_{j=1}^M w_j O_{ij} + \sum_{\substack{p,q \\ p \neq q}} r_{pq} O_{ip} O_{iq} \right). \quad (4)$$

Proposition 6. (Distribution-free guarantee) Consider the class of prediction problems formally characterized in Definition 2. Let also P denote a joint probability distributions over the set of all attribute values $\{O_{ij}\}_{i,j}$, i.e., $\{O_{ij}\}_{i,j} \stackrel{d}{\sim} P$. Then, for any joint probability distribution P the following statement holds true: If $\forall k w_k = v_k$, and $\exists R \in \mathbb{R}$ s.t. $\forall p, q \in \{1, \dots, M\} r_{pq} < R$, and $\exists r > 1 + \theta$ s.t. $\forall k \frac{r-1}{r-(\theta+1)} \binom{M}{2} \theta^2 R \leq v_k \leq \left(\frac{1}{r}\right)v_{k-1}$, then TTB is an optimal strategy for the class of decision-making problems characterized in Definition 2. ■

In simple terms, Proposition 6 formally establishes a distribution-free result granting the optimality of TTB (when generalized to the N -object setting, each with discrete, multi-level attribute values) with respect to the broad class of prediction problems characterized in Definition 2 (with interactions between attributes also accounted for).

Next, Definition 3 formally characterizes a broad class of predictions problems under the *noisy-world* setting wherein the noise component contaminating the prediction problem has a particular structured form: Gaussian distribution.

Definition 3. (Objective function) Let $\forall i \in \{0, 1\}$, O_i denote the two objects a decision-maker should choose from, and $O_{ij} \in \{0, 1\}$ denote the value of the j^{th} attribute of object O_i where $1 \leq j \leq M$. Additionally, let w_k denote the weight corresponding to the k^{th} attribute, and C_i denote the score the object O_i attains on the criterion of interest to the prediction task (e.g., the population of a city, if the prediction task is to predict which of two cities has a higher population). Finally, let $\chi(\cdot)$ be an *arbitrary* monotonically-increasing function (i.e., $\forall x: \frac{d}{dx}\chi(x) > 0$). Then, consider the class of prediction problems satisfying the following:

$$C_i \triangleq \chi \left(\sum_{j=1}^M w_j O_{ij} \right) + \varepsilon, \quad \varepsilon \stackrel{d}{\sim} \mathcal{N}(0, \sigma^2). \quad (5)$$

Proposition 7. (Noise-level-independent probabilistic guarantee) Consider the class of prediction problems formally characterized in Definition 3. Then, for any noise variance $\sigma^2 > 0$, the following statement holds true: If $\forall k w_k = v_k$, and $\exists r \in \mathbb{R}^{>2}$ s.t. $\forall k v_k \leq \left(\frac{1}{r}\right)v_{k-1}$, then the probability with which TTB correctly selects the superior object is ≥ 0.5 .

Proposition 7 formally establishes the following important result for the inherently-noisy world characterized in Definition 4: For any noise level σ^2 , TTB dominates the *selection-purely-by-chance* strategy which select the winning object uniformly at random. This result importantly demonstrates that, independent of noise level σ^2 , the adoption of TTB (instead of the selection-purely-by-chance strategy) is rationally justified for the inherently-noisy, class of prediction problems formally characterized in Definition 3.

Definition 4 below formally characterizes a broad class of predictions problems, once again, under the noisy-world setting; this time, however, the noise component contaminating the prediction problem has an *unstructured* form.

Definition 4. (Probabilistic guarantee) Let $\forall i \in \{0, 1\}$, O_i denote the two objects a decision-maker should choose from. Let also $O_{ij} \in \{0, 1\}$ denote the value of the j^{th} attribute of object O_i , w_k denote the weight corresponding to the k^{th} attribute, and C_i denote the score the object O_i attains on the criterion of interest to the prediction task (e.g., the population of a city, if the prediction task is to predict which of two cities has a higher population). Finally, let $\phi(\cdot)$ be a monotonically-increasing function (i.e., $\forall x: \frac{d}{dx}\phi(x) > 0$). Then, consider the class of prediction problems satisfying the following:

$$\mathbb{P}(C_1 > C_2 | \phi \left(\sum_{j=1}^M w_j O_{1j} \right) > \phi \left(\sum_{j=1}^M w_j O_{2j} \right)) \geq 1 - \eta, \quad 0 < \eta \ll 1.$$

Proposition 8. (Distribution-free guarantee) Consider the class of prediction problems formally characterized in Definition 4. Let also P denote a joint probability distributions over the set of all attribute values $\{O_{ij}\}_{i,j}$, i.e., $\{O_{ij}\}_{i,j} \stackrel{d}{\sim} P$. Then, for any joint probability distribution P the following statement holds true: If $\forall k w_k = v_k$, and $\exists r \in \mathbb{R}^{>2}$ s.t. $\forall k v_k \leq \left(\frac{1}{r}\right)v_{k-1}$, then the probability with which TTB mistakenly selects the inferior object is less than η , where $0 < \eta \ll 1$. ■

Proposition 8 establishes a distribution-free condition sufficient to grant that the probability of TTB erring in a prediction task belonging to the class of problems characterized in Definition 4 is minuscule.

4 General Discussion

In this work, we presented a research program, dubbed formal science of heuristics (FSH), that nicely complements the influential ecological rationality research program (Todd & Gigerenzer, 2007), developing it into a much analytically-richer scientific endeavor. By pursuing its two stated goals (see Introduction section), FSH seeks to (i) mathematically

delineate the key premise ecological rationality rests on—that heuristics are well-matched to the environments in which they are adopted (Todd & Gigerenzer, 2007)—and (ii) establish the strongest analytical results supporting this premise. After all, to rigorously and thoroughly answer whether a heuristic is well-matched to its environment, we need to formally characterize the *broadest* class of environments for which that heuristic performs (near) optimally, and experimentally investigate how often people use that heuristic in such environments.

Instantiating FSH with the well-known Take The Best (TTB) heuristic, and contrary to past analytical work, in this work we considered a much broader class of prediction problems involving nonlinear objective functions (Definitions 1-4) with interactions between attributes also being accounted for (Definition 2). For the classes discussed above, we formally established conditions granting the optimality of TTB when dealing with both non-binary, discrete attribute values (Propositions 3, 5, and 6) and continuous attribute values (Proposition 4). We also analytically investigated a broad class of prediction problems—involving both structured (Definition 3) and unstructured noise (Definition 4)—for which only *probabilistic* guarantees can be provided. Additionally, and in sharp contrast to past analytical work, we also provided distribution-free guarantees on TTB for several classes of prediction problems (Propositions 5–6, and 8).

Our work also serves as a potential template for how FSH could be pursued: For a given heuristic, formally characterize the class of decision-making problems with respect to which the performance of the heuristic is to be analytically investigated, followed by analytical results rigorously delineating the extent to which that heuristic is performing (near) optimally for that class. A generic approach would be to start with a narrow class (containing a set of restricted problems) for which a heuristic is performing (near) optimally; and then gradually expand that class into a larger one and see if previously established performance guarantees still hold (or to establish new performance guarantees, in case they fail to hold). A similar approach has been widely and productively used in theoretical computer science and computational complexity theory, e.g., through formally introducing many complexity classes, with one class serving as a relaxation of another.

Our particular focus on TTB in this work was only meant to showcase how the mindset advocated by FSH could be pursued in the case of a given heuristic—in our case, the Take The Best (TTB) heuristic. Ultimately, a serious investigation of FSH should lead to having mathematically rigorous answers to the two stated goals of FSH for *every* experimentally well-documented heuristic that people use, e.g., the Tallying heuristic (Gigerenzer & Gaissmaier, 2011), the Priority heuristic (Katsikopoulos & Gigerenzer, 2008), the Recognition heuristic (Gigerenzer & Gaissmaier, 2011), and the Minimalist heuristic (Gigerenzer et al., 2008). By now, a large number of heuristics are documented in the literature, many of which still lack an adequate characterization of the en-

vironmental conditions under which they are (near) optimal and/or how deviations from those conditions would lead to performance degradation. Thus, future work following FSH should address this analytical shortcoming.

Rieskamp and Otto (2006) show that people are sensitive to the distribution of cues in an environment, appropriately applying either TTB or a weighted additive mechanism, depending on which will be more accurate. However, how people are able to determine which type of environment they are in has largely remained an open question. Establishing distribution-free guarantees, as advocated by FSH, sheds new light on this open question, by formally demonstrating that a heuristic may well yield adequate performance despite the decision-maker's possibly incomplete (or, in the worst case, erroneous) assumptions about her environmental conditions, thereby liberating her from having a thorough understanding of her environment—a more psychologically plausible assumption. For example, Proposition 5 establishes a condition granting the optimality of TTB (with respect to the class of problems characterized in Definition 1) that holds true for *any* joint probability distribution \mathbb{P} on the set of attribute values. This result has an important implication: Even if the decision-maker makes wrong assumptions about the true underlying distribution \mathbb{P} governing the set of attribute values, adopting TTB still remains to be the optimal strategy (for the class of problems characterized in Definition 1). Crucially, the latter statement remains valid regardless of how wrong the decision-maker's assumptions about \mathbb{P} are.

We must note, however, that the present work (and pursuit of FSH, in general) does not address the recent conundrum raised by Otworowska et al. (2018) regarding the computational intractability of the Adaptive Toolbox theory. As Otworowska et al. (2018) analytically demonstrate, there exists no efficient (i.e., polynomial-time) process that can adapt toolboxes to be ecologically rational for *all* possible environments. A resolution of this complexity-theoretic conundrum might be attained by restricting the class of environments under consideration, based on the psychologically plausible assumption that the range of environments humans have to deal with is undoubtedly vast, but *not* arbitrary.

Pursuit of FSH would have important implications for experimental work on heuristics: Every analytical result (however general it may be) is established under a particular set of assumptions the validity of which needs to be experimentally confirmed. Experimental work should therefore investigate the empirical validity of such assumptions. Likewise, an empirical disconfirmation of an assumption on which an analytical result rests should call for the development of new empirically-grounded formal results. Accordingly, pursuit of FSH yields new experimental work, and, conversely, those experimental findings guide the development of new analytical results—a synergetic scientific endeavor.

Finally, pursuit of FSH allows mathematicians and theoretical computer scientists to make important contributions to the science of heuristics by developing a mathematically-

rigorous foundation for the effectiveness of heuristics in everyday life decisions. As such, we hope that FSH paves the way for having a highly interdisciplinary research program on heuristics wherein analytical and experimental studies, hand in hand, deepen our understanding of the effectiveness of the heuristics we live by. We see our work as a step toward that.

Investigations into human judgment and decision-making have led to the discovery of a multitude of cognitive biases and fallacies, with new ones continually emerging, leading to a state of affairs which can be characterized as the cognitive fallacy zoo! Recently, we have formally presented a principled way to bring order to this zoo (Nobandegani, Campoli, & Shultz, 2019). The work presented here, together with recent formal advances on bringing systematic order to the cognitive fallacy zoo (Nobandegani, Campoli, & Shultz, 2019), suggest a fresh formal approach to pursuing the heuristics-and-biases research program: an approach which aims to lay the formal foundations of the “unreasonable” effectiveness of the heuristics we live by, and to bring mathematically-rigorous systematic order to the cognitive biases ensued by those heuristics.

References

- Baucells, M., Carrasco, J. A., & Hogarth, R. M. (2008). Cumulative dominance and heuristic performance in binary multiattribute choice. *Operations Research*, *56*(5), 1289–1304.
- 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.
- Bröder, A. (2003). Decision making with the “adaptive toolbox”: influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(4), 611.
- Bröder, A., & Gaissmaier, W. (2007). Sequential processing of cues in memory-based multiattribute decisions. *Psychonomic Bulletin & Review*, *14*(5), 895–900.
- Bröder, A., & Schiffer, S. (2003). Take the best versus simultaneous feature matching: Probabilistic inferences from memory and effects of representation format. *Journal of Experimental Psychology: General*, *132*(2), 277.
- Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (2003). Fast, frugal, and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes*, *90*(1), 63–86.
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In *Simple Heuristics that Make us Smart* (pp. 97–118). Oxford University Press.
- Evans, J. S. B., & Over, D. E. (2010). Heuristic thinking and human intelligence: a commentary on marewski, gaissmaier and gigerenzer. *Cognitive Processing*, *11*(2), 171–175.
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, *62*, 451–482.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, *103*(4), 650.
- Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: a brunswikian theory of confidence. *Psychological Review*, *98*(4), 506.
- Gigerenzer, G., Martignon, L., Hoffrage, U., Rieskamp, J., Czerlinski, J., & Goldstein, D. G. (2008). One-reason decision making. In: *Handbook of Experimental Economics Results*. In C. Plott & V. Smith (Eds.), (1st ed., Vol. 1, pp. 1004–1017). North Holland.
- Gigerenzer, G., & Selten, R. (2002). *Bounded rationality: The adaptive toolbox*. MIT press.
- Gigerenzer, G., & Todd, P. M. (1999). *Simple Heuristics that Make us Smart*. New York: Oxford University Press.
- Hoffart, J. C., Rieskamp, J., & Dutilh, G. (2018). How environmental regularities affect people’s information search in probability judgments from experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Hoffrage, U., & Reimer, T. (2004). Models of bounded rationality: The approach of fast and frugal heuristics. *Management Review*, *437–459*.
- Hogarth, R. M., & Karelaia, N. (2005). Simple models for multi-attribute choice with many alternatives: When it does and does not pay to face trade-offs with binary attributes. *Management Science*, *51*(12), 1860–1872.
- Hogarth, R. M., & Karelaia, N. (2006). take-the-best and other simple strategies: Why and when they work well with binary cues. *Theory and Decision*, *61*(3), 205–249.
- Katsikopoulos, K. V., & Gigerenzer, G. (2008). One-reason decision-making: Modeling violations of expected utility theory. *Journal of Risk and Uncertainty*, *37*(1), 35.
- Kearns, M. J., Vazirani, U. V., & Vazirani, U. (1994). *An introduction to computational learning theory*. MIT press.
- Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision*, *52*(1), 29–71.
- Mousavi, S., Gigerenzer, G., & Kheirandish, R. (2016). Rethinking behavioral economics through fast-and-frugal heuristics. *Routledge Handbook of Behavioral Economics*, 280–296.
- Nobandegani, A. S., Campoli, W., & Shultz, T. R. (2019). Bringing order to the cognitive fallacy zoo. In *Proceedings of the 17th International Conference on Cognitive Modeling*. Montreal, QC.
- Nosofsky, R. M., & Bergert, F. B. (2007). Limitations of exemplar models of multi-attribute probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(6), 999.
- Otworowska, M., Blokpoel, M., Sweers, M., Wareham, T., & van Rooij, I. (2018). Demons of ecological rationality. *Cognitive Science*, *42*(3), 1057–1066.
- Pachur, T., Todd, P. M., Gigerenzer, G., Schooler, L., & Goldstein, D. G. (2011). The recognition heuristic: A review of theory and tests. *Frontiers in Psychology*, *2*, 1417.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1988). Adaptive strategy selection in decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(3), 534.
- Payne, J. W., R., B. J., & J., J. E. (1994). *The adaptive decision maker*. New York: Cambridge University Press.
- Rieskamp, J., & Otto, P. E. (2006). Ssl: a theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*(2), 207.
- Shah, A. K., & Oppenheimer, D. M. (2008). Heuristics made easy: An effort-reduction framework. *Psychological Bulletin*, *134*(2), 207.
- Todd, P. M., & Gigerenzer, G. (2000). Précis of simple heuristics that make us smart. *Behavioral and Brain Sciences*, *23*(5), 727–741.
- Todd, P. M., & Gigerenzer, G. (2007). Environments that make us smart: Ecological rationality. *Current directions in psychological science*, *16*(3), 167–171.
- Tversky, A. (1969). Intransitivity of preferences. *Psychological Review*, *76*(1), 31.
- Valiant, L. G. (1984). A theory of the learnable. *Communications of the ACM*, *27*(11), 1134–1142.