

Adaptively Rational Learning

Sarah Wellen¹ · David Danks¹

Received: 15 May 2013 / Accepted: 16 April 2015 / Published online: 24 April 2015
© Springer Science+Business Media Dordrecht 2015

Abstract Research on adaptive rationality has focused principally on inference, judgment, and decision-making that lead to behaviors and actions. These processes typically require cognitive representations as input, and these representations must presumably be acquired via learning. Nonetheless, there has been little work on the nature of, and justification for, adaptively rational learning processes. In this paper, we argue that there are strong reasons to believe that some learning is adaptively rational in the same way as judgment and decision-making. Indeed, overall adaptive rationality can only properly be assessed for *pairs* of learning and decision processes. We thus present a formal framework for modeling such pairs of cognitive processes, and thereby assessing their adaptive rationality relative to the environment and the agent’s goals. We then use this high-level formal framework on specific cases by (a) demonstrating how natural formal constraints on decision-making can lead to substantive predictions about adaptively rational learning and representation; and (b) characterizing adaptively rational learning for fast-and-frugal one-reason decision-making.

Keywords Adaptive rationality · Bounded rationality · Ecological rationality · Heuristics · Learning · Representation · Cognition

Introduction

A core intuition of the bounded/adaptive rationality research programme is that cognitively limited individuals can nonetheless function rationally in an uncertain world by being appropriately “tuned” to their environment. That is, one way to act

✉ Sarah Wellen
sarahadele12@gmail.com

¹ Department of Philosophy, Carnegie Mellon University, Baker Hall 135, Pittsburgh, PA 15213-3890, USA

appropriately in the face of significant constraints is to use processes that are not necessarily reliable in all possible worlds, but are typically reliable in *our* world. The largest strand of this programme has argued that success is due to ‘fast and frugal’ heuristics that take advantage of both the structure of the environment, and also cognitive capacities and limitations of the organism, whether evolved or learned, that are rapid and require relatively little cognitive effort. There has been substantial debate about both the content and suitability of this conception of rationality (e.g., Gigerenzer 1996; Gigerenzer et al. 1999; Kahneman and Tversky 1996; Samuels et al. 2002; Simon 1976), and while these foundational debates are deeply important (we discuss them briefly in “[The Nature of, and Justification for, Adaptively Rational Learning](#)” section), our focus is primarily on questions that are internal to the framework. We largely take for granted that these heuristics are plausible and interesting, and ask about a particular, under-explored type.

Many heuristics have been proposed in recent years for search, inference, judgment, and decision-making (e.g., Gigerenzer and Goldstein 1996; Gigerenzer et al. 1999; Goldstein and Gigerenzer 2002; Hertwig et al. 2012). A well-known example is Take The Best (TTB), a heuristic for multi-attribute decision-making problems (e.g. choosing which of two cities has the greater population based on binary cues such as ‘being the state capitol’ or ‘having a university’). An agent makes a decision using TTB by searching through the available cues in order of their validity and choosing the option that is supported by the first (most valid) cue that discriminates between them. TTB is often supplemented with the Recognition Heuristic, which says that the first cue considered should always be whether the agent even recognizes an option (Goldstein and Gigerenzer 2002). Heuristics such as these typically assume that the agent has cognitive representations that encode relevant aspects of the environment. For instance, TTB (at least in its initial formulation in Gigerenzer and Goldstein 1996) assumed that agents have a representation of the validity of each cue for the specific criterion in question. These representations are critical for the proper functioning of the heuristic, and presumably result from the cognitive agent learning about her environment from her past experiences and reasoning. But apart from some notable exceptions (e.g., Todd and Dieckmann 2005; Hertwig et al. 2003; Hills and Hertwig 2010; Hoffrage et al. 2000; Rieskamp and Otto 2006; Schooler and Anderson 1997), these learning processes are rarely incorporated into models of the heuristics themselves.

In this paper, we argue that learning deserves the same focus and attention from the bounded rationality programme as judgment and decision-making. Agents face many of the same constraints and opportunities while learning as they do while making decisions and inferences. In particular, learning by cognitively limited agents can be significantly improved when it is appropriately tuned to the environment and uses the evolved and acquired capacities of the cognitive agent. There are close interactions between learning and inference/decision-making: the rationality of either depends in part on the other. Many claims that some particular inference or decision-making heuristic is adaptively rational depend implicitly on assumptions about the learning that occurred previously. For instance, the claim that TTB leads to good decision-making in environments with non-compensatory cue structure (i.e. where high validity cues cannot be overridden by several cues with

lower validities) depends implicitly on the assumption that the agent has correctly learned the cue orders (Dieckmann and Todd 2012; Martignon and Hoffrage 1999; but see Katsikopoulos et al. 2010). At the same time, the effectiveness or rationality of particular *learning* processes depend in part on how the representations are used in future judgment and decision-making; a learning process that may be adaptively rational when used by TTB may be inadequate for other decision-making processes.

There are many issues surrounding adaptively rational learning, and this paper can only scratch the surface. We certainly do not intend anything in this paper to be the final word, as research on adaptively rational learning is only just beginning. To help frame the issues appropriately, “[The Nature of, and Justification for, Adaptively Rational Learning](#)” section characterizes what adaptively rational learning even is, and why cognitive agents presumably employ it. “[A General Framework for Adaptively Rational Learning](#)” section then argues that there are certain general features required for any model of adaptively rational learning. We then build a very high-level mathematical framework that includes those features. One key aspect of this framework is that we must make precise what “success” means for a learning heuristic in a given environment, and we argue that this depends critically on the pragmatic and epistemic goals of the learner. “[Adaptive Rationality in the Particular](#)” aims to make this high-level framework concrete by providing both general and particular ways to apply it for particular psychological cases.

The Nature of, and Justification for, Adaptively Rational Learning

The success of a decision-making heuristic depends critically on the availability of the right sort of information. It is common practice to ask how a decision-making heuristic operates in an environment, but we suggest that the question should be worded slightly differently, as the agent does not have direct access to the environment. We should instead ask how a decision-making heuristic operates *for a learner* in an environment, where the information available to the heuristic must be learned or constructed based on limited evidence, time, and computational ability. This more complex question has sometimes been raised for particular domains, as in the course of research on TTB. The initial formulation of TTB (e.g. Gigerenzer and Goldstein 1996) assumed that cue search was based on an accurate representation of the ecological validities of the cues, but controversy soon arose over whether people actually can and do learn cue validities (see, for example, Juslin and Persson 2002; Newell et al. 2004), and if not, whether TTB is really as fast, frugal, and ecologically rational as it initially appeared. That is, one objection to TTB was that no adaptively rational learning counterpart had been provided for it.

Several follow-up studies investigated how participants perform when they must learn about the cues on their own during decision-making and found that neither their search orders nor their ratings of cue “usefulness” matched the ecological validities that TTB assumed as input (Newell et al. 2004; Dieckmann and Todd 2012). More generally, some researchers contended that the computational and memory requirements for learning cue validities are too complex for bounded

agents (Juslin and Persson 2002; Newell et al. 2004; Dougherty et al. 2008). Gigerenzer et al. (1999) had initially proposed that people could learn the cue validities by tracking the frequency of correct responses among the cases in which the cue discriminates, but Newell et al. (2004) demonstrated that this requires a fairly complex set of computations. For each cue, the learner has to track the frequency of correct discriminations as well as the frequency of cases in which the cue discriminates, and then compute a percentage from these two frequencies. Newell, et al. proposed instead that people use the “success rate,” which measures the proportion of correct decisions that would be made by using a cue on its own. Success rate incorporates both the validity and ‘discrimination rate’ of a cue, and thus penalizes cues that rarely help the agent choose between two options. More importantly for our purposes, the success rate *ordering* (which is all that is necessary for determining search order) can be learned by monitoring only the frequency of correct discriminations for each cue. Simulations of TTB with various search orders have revealed that cue validity is only one of many orders that produces good accuracy on the task (Martignon and Hoffrage 1999; Katsikopoulos et al. 2010). A further set of simulations (Todd and Dieckmann 2005; Dieckmann and Todd 2012) focused on simple heuristics for learning cue search orders ‘online’ during decision making, and demonstrated that simple “swap” and “tally” rules can result in remarkably good performance despite their computational simplicity.

We thus have an example of how learning could be approached from the adaptive rationality framework; this section asks whether we can generalize from it to broader morals. As this example suggests, learning can matter in two different ways for adaptive rationality. First, the learning process determines the information—type, scope, and representation—that is actually available to the decision-maker in the relevant environments. Decision-making methods almost always require certain types of information to be provided as input, and so the actual cognitive representations produced by the learning process help determine whether certain decision-making methods are even possible for a cognitive agent. More generally, an evaluation of the rationality of a decision making process is only possible if we know which features of the environment are represented at the time of judgment or decision. Second, the extent to which the learning mechanism is itself fast and frugal matters for the agent’s overall adaptive rationality. Adaptive rationality is a property of cognitive agents more generally, not just particular cognitive processes. If a decision-making heuristic only performs well when the learner has exhaustive evidence and performs extensive computations, then this heuristic cannot contribute to fast and frugal cognition regardless of how it appears on its own. Adaptive rationality requires not only that inference and decision-making be fast and frugal given some representations, but also that those representations be acquired through fast and frugal learning.

These connections are arguably unsurprising since learning itself involves certain types of decision-making, such as “deciding” which cognitive representation to encode given perceptual evidence. These decisions are quite different than those usually studied in the adaptive rationality programme, however, as they do not directly manifest in observable behavior. And while the importance of learning has been studied in particular cases (e.g., for the performance of TTB), we arguably lack

a general framework for analyzing what makes learning adaptively rational. We suggest that learning methods are not free from the constraints and limits that plague other types of inference and decision-making, and so all of the arguments in favor of adaptively rational inference and decision-making apply equally well to the case of learning. We here consider three arguments that notions of adaptive rationality could perhaps provide a fruitful lens through which to understand how we learn cognitive representations.

First, as with decision-making, certain standard assumptions of axiomatically rational¹ learning methods are violated in many real-world learning situations, yet people are often able to learn in sensible ways in these cases. In particular, essentially all axiomatically rational learning methods assume that there are both a fixed and known space of possibilities (a “hypothesis space”), and also known and computable functions connecting those possibilities with the evidence that the learner might see. Real-world learning, however, typically requires that one leaves the door open to the possibility of “something else” that has not previously been considered or specified, and so the possibility space cannot necessarily be specified in advance of learning. Moreover, even if we have well-specified possibilities, we often lack the relevant information to fully specify or compute the implications of particular possible hypotheses. Although axiomatically rational learning methods have appealing formal and mathematical properties, they gain their justification only by virtue of provable properties, not probable or possible ones. Real-world learning rarely provides the resources to obtain such guarantees, and so we should arguably consider other ways to analyze the rationality of learning.

A second argument for focusing on adaptively rational learning methods arises from general observations involving overfitting and the well-known bias/variance tradeoff (Geman et al. 1992; see Gigerenzer and Brighton 2009 for connections with cognitive processes). A presumably necessary condition for the axiomatic rationality of any learning method is that it yields the correct answer, if possible, given sufficient data. Put more precisely, axiomatically rational learning methods must be *unbiased*, as any bias will lead to errors in the long-run. In addition, the possibility space for a learning method must include as much complexity as in any possible learning situation, including those in which we receive arbitrarily large amounts of data. These two requirements jointly imply that axiomatically rational learning methods will typically exhibit substantial variance when provided only limited evidence. That is, such methods will almost always be overly sensitive to chance variations in the input data: different (small) random samples from the same “ground truth” can yield wildly different outputs for the learning method. Real-world learning situations are frequently characterized by limited evidence, sometimes only a single datapoint. It can thus sometimes be quite damaging to the learner to use an axiomatically rational learning method, precisely because the high variance of such methods on small evidence sets can translate into large numbers of errors. Instead, it is often in the learner’s best interests to use a method that is perhaps biased, but also has substantially reduced variance. The question of

¹ That is, methods for which there are proofs of their asymptotic and/or short-run reliability.

where to stand on this bias-variance tradeoff in learning can only be asked, however, if we adopt the standards of adaptive, rather than axiomatic, rationality.

This last observation points towards the third argument, which suggests that the “proper” standard of rationality is external (e.g., successful actions and outcomes) rather than internal (e.g., valid reasoning), and so adaptive rationality is the proper way to evaluate all cognitive processes, including learning (see discussions in Gigerenzer et al. 1999; Samuels et al. 2002; Simon 1976). There is a voluminous literature about arguments of this type (Kahneman and Tversky 1996; Gigerenzer 1996) as well as attempts to better understand the different types of rationality (e.g., Hammond 1996), and we do not wish to get bogged down in those issues. Nonetheless, we suggest that these arguments are, while not completely unproblematic, sufficiently plausible that we should see how far an investigation into adaptively rational learning can take us.

Of course, one might object that learning does not directly affect the external world in the same way as decision-making, and so it cannot be evaluated according to an external standard of rationality. Learning might be thought to interface with the world only indirectly: the output of learning is not an action that is successful or unsuccessful, but rather a representation that is inefficacious on its own. That is, one might object that there is no “external manifestation” on which to evaluate the adaptive rationality of learning.

This objection raises serious concerns, but complementary ones arise equally for inference and decision-making: those cognitive processes require cognitive representations as input, and so they seem to lack anything appropriately “external” on the *input* side. In general, any external standard of rationality should properly apply to pairs of learning and inference/decision processes. In practice, we (as cognitive scientists) typically fix one method or the other as “obvious” and then focus on the other member of the pair, but this is a substantive choice. For any particular learning method, there will be inference/decision methods for which the resulting pair yields externally irrational behavior, and similarly for any particular inference/decision method. Thus, to the extent that one endorses arguments in favor of an external standard of rationality, one must accept that this standard applies equally well to learning. Decision-making may mediate the effect of learning on actions and outcomes, but it does not thereby eliminate learning’s impact. External rationality applies to learning-inference/decision pairs if to anything at all, and so we must consider whether learning also employs processes that embody adaptive rationality. Of course, all of these arguments leave open the exact form and content of adaptively rational learning, and so we now turn to better understanding it.

A General Framework for Adaptively Rational Learning

A learning process will be adaptively rational to the extent that bounded individuals can use it to achieve their goals in appropriate environments with limited and uncertain evidence. As just noted, this assessment depends in part on the accompanying decision-making process; for now, we assume a fixed judgment or decision-making process and consider the “choice” of a relevant process later.

Learning and decision-making involve many of the same challenges (e.g., limited and uncertain evidence, time constraints), so we propose that agents employ learning heuristics with the same qualitative features as the heuristics found in search, inference, and decision-making. First, the processes should be fast and frugal; that is, they should not require too much evidence or computation in order to draw useful conclusions. There is significant ambiguity about exactly how to measure evidential and computational costs, but we leave those issues aside.

Second, adaptively rational learning should take advantage of evolved and acquired capacities. In some cases, these capacities may be the same as those proposed for decision-making heuristics (e.g., counting, frequency monitoring, imitation), though learning may also employ capacities that have not previously been discussed. In particular, perceptual and attentional processes will presumably be important for determining what evidence is available to the agent at a given time. It is also possible that relatively automatic and unconscious learning processes (e.g., sequence learning, reward learning, association formation, recognition) can provide input to more complex inferential processes. Although many papers in the adaptive rationality programme focus on evolved capacities, we are deliberately agnostic about whether the relevant cognitive capacities have a phylogenetic (i.e., evolutionary) or ontogenetic (i.e., individual learning) basis. Organisms can certainly learn how to learn, and so we see no reason to assume that only evolved capacities are relevant. In our view, what matters for adaptively rational learning is that it involves the reuse of capacities for multiple purposes, rather than assuming that there is an entirely new learning capacity for each learning challenge.

The final criterion is that the heuristic must be ecologically rational; that is, it must allow the agent to achieve her goals in her environment. As discussed above, this poses a significant conceptual challenge because the outputs of a learning process do not naturally interface with the environment in the same way as the outputs of a decision process. One response is to assume that the ecological rationality of learning is measured by its ability to arrive at *true conclusions* in the relevant environments. Much of the learning literature implicitly assumes this standard, and heuristics have been defended using this criterion. For instance, people appear to use order of observations as a cue to causal structure (Lagnado et al. 2007), even when they know that, in the artificial setting of the psychology lab, observational order may not follow the true sequence of events (Lagnado and Sloman 2006; Sloman and Lagnado 2005). This ‘temporal sequence’ heuristic plausibly takes advantage of our evolved capacity to detect temporal sequences, draws conclusions based on only a small set of observations, and usually produces accurate mental representations in environments (like ours) where observations of causes typically precede observations of effects. Another example is Monte Carlo approximation of Bayesian inference (e.g., Denison et al. 2013; Shi et al. 2010; Bonawitz et al. 2011), which may show how a learning process could typically produce correct conclusions by approximating (in a bounded way) complex functions using simple evolved capacities (e.g., exemplar-based reasoning).

We contend, however, that defining ecological rationality as the propensity to arrive at true conclusions about the environment is an error, as it presupposes that agents have only epistemic goals for learning (if those are even possible; Danks, in

press). Cognitive agents will almost always have at least some pragmatic goals: to achieve the best outcome, receive as much money as possible, appear competent to others, and so forth. Moreover, there are various psychological results that suggest that these pragmatic goals can have a significant impact on the cognitive representations that people learn, such as in concept learning (e.g., Bailenson et al. 2002; Chin-Parker and Ross 2002; Hoffman and Rehder 2010; Markman and Ross 2003; Ross 1997, 2000; see also Danks 2014). The adaptive rationality perspective provides the conceptual resources to incorporate these pragmatic goals. In particular, learning for pragmatic goals can be successful even if it results in incomplete or inaccurate representations of the world; the truth is not always best, or even necessary, for success at a task. Of course, randomly generated cognitive representations typically will not lead to success either, but the adaptive rationality framework allows us to begin to ask these questions in a principled way.

We now turn to the challenge of mathematically representing adaptively rational learning. We recognize that there might be alternative ways to capture adaptively rational learning formally, but the precise mathematical details are less important than the qualitative pieces that are required.

In general, the success of learning is determined by whether the acquired representations allow the agent to achieve her goals—epistemic *and* pragmatic—in subsequent inference and decision-making. The overall picture that we adopt can be expressed graphically as: *World* \rightarrow *Evidence* \rightarrow *Cognitive representation* \rightarrow *Action*. That is, the world generates evidence from which the agent generates a cognitive representation (via learning) that subserves action (via inference and decision-making). In this model, the *Evidence* \rightarrow *Cognitive Representation* link denotes the learning method L ; and *Cognitive Representation* \rightarrow *Action* refers to the inference or decision process D . This simplistic picture can be complicated in many different ways, such as including an *Action* \rightarrow *Evidence* connection to capture the ways in which our decisions shape the evidence that we see in the future. Nonetheless, this simple picture is already sufficient to reveal some of the complexities of adaptively rational learning, such as evaluating the adaptive rationality of actions given the world or evidence by considering $\langle L, D \rangle$ pairs, rather than either one in isolation. Finally, we must include a representation of an agent's goal G , which we operationalize through a value function $V_G(a, w)$ that depends on the particular action a and world w . For example, if the goal is simply to learn some feature of the world, then the value function would be maximal when a “matches” w in the relevant respects.

Given this general setup, we can mathematically evaluate any particular $\langle L, D \rangle$ pair of learning and decision processes. Learning given evidence e will yield cognitive representations that are used by the decision process, along with knowledge of the goal value function, to produce an action. Moreover, these different learning and decision processes are presumably all probabilistic: $P(a|e) = D(L(e), V_G)$. Both L and D can be generalized in obvious ways to include background knowledge; in particular, L can depend on the goal value function, if it is known prior to learning. For a fixed set of goals \mathbf{G} and probability distributions over those goals and the world-states, the expected value of a particular $\langle L, D \rangle$ pair can be expressed as:

$$EV(\langle L, D \rangle) = \sum_{G \in \mathbf{G}} \sum_{w \in \mathbf{W}} \sum_{e \in \mathbf{E}} \sum_{a \in \mathbf{A}} V_G(a, w) D(L(e), v_G) P_{\text{sample}}(e|w) P(w) P(G)$$

We previously mentioned cognitive representations, but they do not explicitly appear in this equation. In this particular formulation, they play only an implicit role: the space of cognitive representations determines the amount of information (about e) that can possibly be transmitted through $L(e)$ to the decision process D . As we consider more specific situations, they will figure more prominently.

The expected value equation given above provides a general standard by which to evaluate the adaptive rationality of particular $\langle L, D \rangle$ pairs. A wide range of current learning and decision methods can be represented in it with minimal adjustments or modification. At the same time, this equation is arguably *too* general: it is rather opaque, and it is hard to see how to draw any general conclusions from it. We thus consider several more specific cases in the next section to show both how to evaluate the adaptive rationality of learning methods, and also the distinctive predictions of such an analysis.

Adaptive Rationality in the Particular

This section explores both general formal constraints and substantive situation-specific ones that help make this framework more concrete and usable. The first half of this section is more mathematical in nature, the second half more psychological. For the mathematical investigation, we first explicitly model the space of possible cognitive representations \mathbf{R} . A particular cognitive representation r can be understood to encode a “way that the world could be.” Moreover, we assume that there is typically a determinate best (for a goal) action for any particular state of the world. There may be unusual circumstances in which indeterministic choice is best even when the exact state of the world is known (e.g., the use of a truly mixed strategy in a game), but most cases of uncertainty about actions arise precisely because we do not actually know how the world is. These observations suggest a natural constraint on the decision function D , if the cognitive agent’s uncertainty is expressed by probabilities over cognitive representations (i.e., $L(e) = P(\mathbf{R})$):

Determinism: Recall $D(P(\mathbf{R}), V_G) = P(\mathbf{A})$. If $P(r) = 1$, then for some $a \in \mathbf{A}$, $D(P(\mathbf{R}), V_g) = 1$. For convenience, we write $D(r, V_G) = a$ to indicate that the cognitive representation r dictates action a (given goal G).

In other words, the decision-procedure dictates a unique best action for every different goal and “way the world could be.” At the same time, we might expect that there are no “interactions” in this uncertainty: the decision probability of an action a should be simply the probability of the cognitive representations that (deterministically) would lead to a , if they were fully believed or endorsed. More precisely:

Composition: $D(P(\mathbf{R}), V_g) = P(\mathbf{A})$ such that $P(a) = P(\{r : D(r, V_g) = a\})$.

Determinism and Composition are natural constraints on the decision function, and many standard decision algorithms from the adaptive rationality framework satisfy these two properties. More interestingly for our focus, a decision function satisfying Determinism and Composition naturally groups together cognitive representations that are “the same” with respect to the decision process. More precisely, if D satisfies Determinism and Composition and $D(r_1, V_G) = D(r_2, V_G)$, then the cognitive agent’s decision will be indifferent over reallocation of probabilities between r_1 and r_2 ; only the sum of those probabilities are relevant for decision-making. For example, if I am choosing between two urns with the goal of getting a black ball, and the only representations with positive probability are $r_1 =$ “Urn 1 has 75 % black balls and Urn 2 has 60 % black balls” and $r_2 =$ “Urn 1 has 70 % black balls and Urn 2 has 50 % black balls,” then the relative probabilities assigned to r_1 and r_2 should not matter; I will choose Urn 1 regardless. Formally, any D that satisfies Determinism and Composition induces a *partition* of the cognitive representations \mathbf{R} —that is, a division of \mathbf{R} into disjoint subsets—where two representations r_i and r_j are in the same partition element if and only if $D(r_i, V_G) = D(r_j, V_G)$. More colloquially, each partition element is a set with just those representations that all lead to the same action, and so are indistinguishable with respect to the decision process.

This partitioning of the cognitive representations yields freedom in an adaptively rational learning process. If D does not distinguish between r_1 and r_2 with respect to goal G , then learning methods that differ only in output about r_1 and r_2 will exhibit the same performance, and so other factors can be relevant in determining their adaptive rationality. More precisely, we can talk about two learning methods L_1 and L_2 being indistinguishable relative to a partition π for a decision process D just when, for each possible piece of evidence, the outputs of L_1 and L_2 are the same when assessed at the level of elements of π .² If an agent’s cognitive representations are at the level of partition elements,³ then she should (on representational grounds) be indifferent between L_1 and L_2 , and so one learning process can be preferred for more pragmatic reasons, such as which method exploits evolved or acquired capacities. This conclusion is at odds with many analyses of the rationality of particular learning processes, at least if the cognitive representations in \mathbf{R} are assumed to be sufficiently fine-grained. For example, most arguments that Bayesian updating provides a rational learning method will not apply to alternative learning methods that are indistinguishable relative to a partition, even though those methods may be equally *adaptively* rational (and perhaps more so if we take into account, for example, the time the brain would need to complete each computation). Moreover, for a particular goal G and decision process D that satisfies Determinism and Composition, the partition induced by D and V_G is naturally privileged and so we see how learning and decision processes start to fit together: the goal and decision tell us which distinctions matter, and learning tries to extract and represent those distinctions from the environment.

² In particular, the sums of probabilities within each partition element must be the same.

³ There are many ways to have “representations at the level of partition elements,” including both coarsened representations and inattention to particular distinctions in the world.

Importantly, this way of thinking about learning provides us with a principled way to determine how one ought—given a goal G and decision process D —to coarse-grain one’s cognitive representations. The representations in \mathbf{R} can potentially be quite fine-grained, and so more specific than is required for either learning or decision. For example, when learning about the two urns, a seemingly natural \mathbf{R} consists of all ordered pairs of numbers between zero and 100, $\langle U_1, U_2 \rangle$, representing the percentages of black balls in Urn 1 and Urn 2. If, however, our goal is to choose an urn so as to pick a black ball, then only three cognitive representations are decision-relevant: “Urn 1 has more black balls than Urn 2”; “Urn 2 has more black balls than Urn 1”; and “Urn 1 has the same number of black balls as Urn 2”. Formally, the natural decision process partitions the ordered pairs of \mathbf{R} into three subsets: (i) those in which $U_1 > U_2$; (ii) those in which $U_1 < U_2$; and (iii) those in which $U_1 = U_2$. There is no need for the agent to include any finer distinctions in her cognitive representations. Many different learning processes will be equally effective for this goal and decision process, and so adaptive rationality may depend deeply on other factors.

The preceding discussion focused on the case of a single, known goal, but we often have multiple goals or uncertainty about our future goals. If there are multiple goals, then the normative π depends partly on whether there is a cost to use partitions π with more elements; in psychological terms, is it costly for the cognitive agent to represent more possibilities? If there is no such cost, then the optimal partition may well be the one in which each element contains exactly one cognitive representation (i.e., nothing gets grouped together).⁴ More generally, though, there might plausibly be a “cost” (of one sort or another) to using a partition with many elements. In that case, the adaptively rational way to coarse-grain one’s cognitive representations will depend on a complicated trade-off between (a) this cost; and (b) having a partition that is sufficiently rich to capture the distinctions that matter for the different goals (with their different likelihoods of occurring). In either case, we can start to see how this approach can lead to interesting predictions about how adaptively rational learners should *ignore* some distinctions in the world, depending on their goals. That is, adaptively rational learners can be normatively justified in ignoring learnable differences in the environment, as long as those differences do not make a difference for foreseeable decisions.

We can also use this general framework to better understand particular proposed heuristics. We focus here on the psychological problem of learning for one-reason decision-making on a binary choice task. In particular, what representations ought we learn and use with the TTB heuristic? In the standard TTB heuristic (Gigerenzer and Goldstein 1996) the cognitive agent searches through cues sequentially in order of their ecological (i.e. population-level) validity (search rule), stops when a cue is found that discriminates between the two items (stop rule), and chooses between the items based solely on the discriminating cue (decision rule). Since the ecological validities are not known but must be learned, we consider a slight modification—

⁴ More precisely, for each representation r , let O_r be the vector of “optimal responses” for the N goals (i.e., $O_r = \langle a_1, \dots, a_N \rangle$ where $a_i = D(r, V_{G_i})$). It is straightforward to prove: if all partitions are equally costly, then r and s are in the same element of the optimal partition if and only if $O_r = O_s$.

call it TTB*—in which the search rule depends on the (uncertain) cognitive representations $P(\mathbf{R})$.⁵ This modification allows us to analyze which features of the cognitive representation are necessary for TTB* to deliver ecological success without making a priori assumptions about the structure of the representations themselves.

TTB* satisfies Determinism and Composition if we add a few small constraints. Recall that a decision procedure satisfies Determinism whenever each possible representation r of the world leads deterministically to a specific action. This can be violated by TTB* if two cues had identical represented usefulness. To avoid this problem, we assume that no fine-grained representation encodes two cues as having *exactly* equal usefulness; alternately, we could incorporate a deterministic tie-breaking mechanism. Importantly, this assumption does *not* imply that the agent will always search the cues in a deterministic order, as the agent may have uncertainty about the exact nature of the world (i.e., the exact, correct cue order). In order to satisfy Composition, we make the minimal assumption that the probability of using cue A (rather than cue B) is based on the agent's representation of the probability that cue A is more useful than cue B. Assuming these two conditions are satisfied, TTB* imposes a partition over the space of possible representations \mathbf{R} : r_i and r_j are in the same partition element if and only if they imply the same usefulness ranking for the cues, though not necessarily the same exact numbers. That is, the only feature of the agent's cognitive representation that affects her ecological success is the *rank order* of the cues, rather than any absolute measures.

Importantly, the preceding analysis did not assume that the fine-grained representations \mathbf{R} must directly represent search orders. Rather, the conclusion that rank order, and so search order, is the only representation that “matters” was derived from the structure of the environment, the agent's goal, and TTB*. The cognitive representations could include information about search order, cue validity, success rate, reward value, positive emotional valence, or any combination thereof. But once we specify those representations \mathbf{R} , we find that qualitative search order is the only feature of $P(\mathbf{R})$ that influences ecological success. This narrow focus can lead to considerable benefit for the agent: in particular, using a more coarse-grained hypothesis space may reduce the variance of learning, though perhaps at the cost of increased bias. As we noted in “[The Nature of, and Justification for, Adaptively Rational Learning](#)” section, this trade-off can often improve accuracy in the short- (or at least not infinitely long-) run. But we should also bear in mind that more precise representations may be useful if we have other goals, or if we use them for decision-making methods other than TTB*. Alternately, more precise representations might actually be easier to learn; for example, precise success values might be computed for other reasons, or be produced by some already-present capacity.

The decision procedure on its own allows us to define the features of the representation that can *possibly* affect ecological success (i.e. the partitions), but the benefit to the agent of assigning credence to one partition over another depends on her goals and environment. The “best” partition may result in much more success than the second-best one, but often they may be almost identical. In this vein,

⁵ \mathbf{R} could be the space of possible cue validities or another similar representational structure.

Martignon and Hoffrage (2002) examined the performance of the possible cue orderings (362,880 in total) for TTB in the German cities environment, and showed that there is a range of orderings that perform as well as (or even better than) the rankings implied by the true cue validities. The true cue validities have a 74.2 % accuracy rate, which falls towards the top of the range (62–75.8 %) but only 4.2 % points above the accuracy achieved by choosing cues at random (70 %). The value function over the partitions is of course environment- and goal-specific, but we can make some general observations that will be true for all particular cases. Whenever multiple cues have very similar usefulness, there is relatively little practical benefit to getting their search order exactly right. Conversely, there is quite a substantial benefit whenever cues vary widely in their utility. As a result, the cue structure of the environment (in combination with the decision procedure) may result in a scenario where adopting a hypothesis space that is even more coarse-grained than the partition (e.g., sorting cues into n usefulness classes where n is less than the number of cues) could further reduce variance and improve ecological success.

The partition determined by the decision procedure TTB* is the maximally fine-grained one for that decision procedure and goal, but even those distinctions might not be worth the cost given other goals, environmental structures, and computational capacities of the agent. More generally, we conjecture that bounded, but adaptively rational, agents should exhibit considerable flexibility in trading accuracy for improved performance on other measures. The classic example is the trade-off between accuracy and the number of cues retrieved (the ‘success rate’ measure explicitly makes this trade-off), but there may be other trade-offs that are beneficial only from a learning perspective. For instance, agents might be frugal in the amount of evidence they use because their learning mechanism uses less information, either by searching fewer cues or by attending more to cues that have not previously been searched. Similarly, a learning mechanism that requires less computation to update the agent’s cognitive representations might naturally be preferred over a learning mechanism that requires more.

Dieckmann and Todd (2012) examined several algorithms for iteratively learning cue orders. For instance, they considered a simple swap rule where a cue is moved up in the search order if it makes a correct discrimination, and down if it makes an incorrect one. This proposed rule fits well with the requirements of adaptive rational learning because (i) only the necessary information—search order—is represented; (ii) it is frugal in the amount of information and complexity of updating rules; and (iii) it produces fairly good performance (71 vs. 74 % for TTB). One particularly promising feature of this rule is that it works by using limited information (i.e., only the outcome of the current choice) to make a local modification to the representation. This feature is characteristic of greedy algorithms, which have been shown in machine learning to often reach locally optimal solutions without extensive memory or computational requirements. This class of algorithms could provide significant inspiration going forward.

Finally, while considerable work has been done on evaluating whether one-reason decision-making is robust across search orders, our framework allows us to ask the converse question: is a given learning mechanism robust across different decision-making strategies? There is considerable evidence that people do not

universally use one-reason decision-making (Newell and Shanks 2003; Bröder 2000, 2003), so one might wonder whether adaptively rational learning mechanisms could produce appropriate representations for these other decision procedures. For example, the success of improper linear models in many decision tasks (e.g., Dana 2008; Dawes 1979; Lovie and Lovie 1986; von Winterfeldt and Edwards 1986) suggests the possibility that learning in these cases is over a greatly coarsened set of representations. More generally, an analysis of learning for other decision-making strategies (e.g., two- or three-reason decision making, tallying, weighted additive models, decision trees) is beyond the scope of this paper, but our model provides the conceptual resources and motivation to begin exploring this uncharted territory.

Conclusion

Our goal in this paper was to identify a significant gap in the adaptive rationality research programme and a significant opportunity for learning research. The adaptive rationality of a cognitive agent depends not just on her inference, judgment, and decision-making strategies, but also on her learning processes. We should expect that people use fast and frugal learning methods that are based on evolved cognitive capacities, and provide the information required for them to reach their goals, both epistemic and pragmatic. This paper has been admittedly programmatic in places, as the nature of adaptively rational learning is clearly a bigger subject than we can exhaustively examine in a single article. Nonetheless, we contend that the formal framework we have presented—both in general, and in the particular—provides the conceptual resources that are required to explore the possibility that human learning is adaptively rational.

Acknowledgments Thanks to participants at the 2013 workshop on Finding Foundations for Bounded and Adaptive Rationality for helpful comments and criticisms, particularly Stefan Herzog, Ralf Mayrhofer, and Peter Todd. Research for this paper was partially supported by a James S. McDonnell Foundation Scholar Award to Danks.

References

- Bailenson, J. N., Shum, M. S., Atran, S., Medin, D. L., & Coley, J. D. (2002). A bird's eye view: Biological categorization and reasoning within and across cultures. *Cognition*, *84*, 1–53.
- Bonawitz, E., Denison, S., Chen, A., Gopnik, G., & Griffiths, T. L. (2011). A simple sequential algorithm for approximating Bayesian inference. *Proceedings of the Thirty-third Cognitive Science Society*.
- 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*, 1332–1346.
- 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*, 611–625.
- Chin-Parker, S., & Ross, B. H. (2002). The effect of category learning on sensitivity to within-category correlations. *Memory and Cognition*, *30*, 353–362.
- Dana, J. (2008). What makes improper linear models tick. In J. I. Krueger (Ed.), *Rationality and social responsibility: Essays in honor of Robyn Mason Dawes*. Mahwah: Lawrence Erlbaum Associates.

- Danks, D. (2014). *Unifying the mind: Cognitive representations as graphical models*. Cambridge: The MIT Press.
- Danks, D. (in press). Goal dependence in (scientific) ontology. *Synthese*.
- Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, *34*, 571–582. doi:10.1037/0003-066X.34.7.571.
- Denison, S., Bonawitz, E. B., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The sampling hypothesis. *Cognition*, *126*(2), 285–300.
- Dieckmann, A., & Todd, P. M. (2012). Simple rules for ordering cues in one-reason decision making. In P. M. Todd & G. Gigerenzer (Eds.), *Ecological rationality: Intelligence in the world* (pp. 274–308). Oxford: Oxford University Press.
- Dougherty, M. R., Franco-Watkins, A. M., & Thomas, R. (2008). Psychological plausibility of the theory of probabilistic mental models and the fast and frugal heuristics. *Psychological Review*, *115*(1), 199–213.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, *4*, 1–58.
- Gigerenzer, G. (1996). On narrow norms and vague heuristics: A rebuttal to Kahneman and Tversky. *Psychological Review*, *103*, 592–596.
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, *1*, 107–143.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*(4), 650–669.
- Gigerenzer, G., Todd, P. M., & the ABC Research Group. (1999). *Simple heuristics that make us smart*. New York: Oxford University Press.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, *109*, 75–90.
- Hammond, K. R. (1996). *Human judgment and social policy: Irreducible uncertainty, inevitable error, unavailable injustice*. New York: Oxford University Press.
- Hertwig, R., Fanselow, C., & Hoffrage, U. (2003). Hindsight bias: How knowledge and heuristics affect our reconstruction of the past. *Memory*, *11*, 357–377.
- Hertwig, R., Hoffrage, U., & the ABC Research Group. (2012). *Simple heuristics in a social world*. Oxford: Oxford University Press.
- Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science*, *21*, 1787–1792.
- Hoffman, A. B., & Rehder, B. (2010). The costs of supervised classification: The effect of learning task on conceptual flexibility. *Journal of Experimental Psychology: General*, *139*(2), 319–340.
- Hoffrage, U., Hertwig, R., & Gigerenzer, G. (2000). Hindsight bias: A by-product of knowledge updating? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 566–581.
- Juslin, P., & Persson, M. (2002). PROBABILITIES from EXemplars (PROBEX): a “lazy” algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, *26*, 563–607.
- Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, *103*, 582–591.
- Katsikopoulos, K. V., Schooler, L. J., & Hertwig, R. (2010). The robust beauty of ordinary information. *Psychological Review*, *117*, 1259–1266.
- Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*, 451–460.
- Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation: Cues to causal structure. In A. Gopnik & L. E. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 154–172). Oxford: Oxford University Press.
- Lovie, A., & Lovie, P. (1986). The flat maximum effect and linear scoring models for prediction. *Journal of Forecasting*, *5*, 159–168. doi:10.1002/for.3980050303.
- Markman, Art B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin*, *129*, 592–613.
- Martignon, L., & Hoffrage, U. (1999). Where and why is Take The Best fast, frugal and fit? In G. Gigerenzer & P. M. Todd (Eds.), *Simple heuristics that make us smart* (pp. 119–140). Oxford: Oxford University Press.
- Martignon, L., & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision*, *52*, 29–71.

- Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of “success”. *Journal of Behavioral Decision Making*, *17*(2), 117–137.
- Newell, B. R., & Shanks, D. R. (2003). Take the best or look at the rest? Factors influencing ‘one-reason’ decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 53–65.
- Rieskamp, J., & Otto, P. E. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, *135*, 207–236.
- Ross, B. H. (1997). The use of categories affects classification. *Journal of Memory and Language*, *37*, 240–267.
- Ross, B. H. (2000). The effects of category use on learned categories. *Memory and Cognition*, *28*(1), 51–63.
- Samuels, R., Stich, S., & Bishop, M. (2002). Ending the rationality wars: How to make disputes about human rationality disappear. In R. Elio (Ed.), *Common sense, reasoning, and rationality* (pp. 236–268). New York: Oxford University Press.
- Schooler, L. J., & Anderson, J. R. (1997). The role of process in the rational analysis of memory. *Cognitive Psychology*, *32*, 219–250.
- Shi, L., Griffiths, T. L., Feldman, N. H., & Sanborn, A. N. (2010). Exemplar models as a mechanism for performing Bayesian inference. *Psychonomic Bulletin & Review*, *17*(4), 443–464.
- Simon, H. A. (1976). From substantive to procedural rationality. In S. J. Latsis (Ed.), *Method and appraisal in economics*. Cambridge: Cambridge University Press.
- Sloman, S. A., & Lagnado, D. A. (2005). Do we “do”? *Cognitive Science*, *29*, 5–39.
- Todd, P. M., & Dieckmann, A. (2005). Heuristics for ordering cue search in decision making. In Lawrence K. Saul, Yair Weiss, & Leon Bottou (Eds.), *Advances in Neural Information Processing Systems 17* (pp. 1393–1400). Cambridge: MIT Press.
- von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research*. Cambridge: Cambridge University Press.