

The Evaluative Mind

16

Julia Haas
2023

16.1 Introduction

In his original introduction to *Mind Design*, John Haugeland observes that “an ‘experiment’ in mind design is more often an effort to *build* something and make it work, than to observe and analyze what already exists” (Chapter 2, this volume). But what happens when such an experiment in mind design succeeds in a way and to a degree that few could have predicted? And more to the point, how do we take on the implications of such successes when they challenge central features of how we understand the mind?

These are the circumstances that we find ourselves in with respect to developments in reinforcement learning. Over the past twenty-five years, reinforcement learning has had a tremendous impact on the development of artificial intelligence and has been a major driver in advancements in the so-called ‘decision sciences’—computational neuroscience, neuroscience, psychology, psychiatry, and economics. But even as we continue to advance the notion of reward maximization as a general solution to the problem of artificial intelligence (Silver, 2015), we have not yet embraced the full implications of reinforcement learning, together with the accompanying reward-prediction hypothesis, for our conceptions of the mind. That is, we continue to think of the mind as some form of a thinking machine (e.g., “thinking, intellect,” Haugeland, Chapter 2, this volume), where such thinking is best understood as some type of computation—ecumenically including neural networks, deep learning, genetic algorithms, and so on.

I propose that the successes and contributions of reinforcement learning urge us to see the mind in a new light, namely, to recognize that the mind is fundamentally *evaluative* in nature. There are weaker and stronger versions of this thesis.

The weaker version, which I commit to here, proposes that the mind is, at a fundamental level, in the business of evaluating states of affairs as better or worse. This version is additive in nature: it says that, *in addition to* performing computations over representations of descriptive matters of fact, the mind *also* performs computations over representations of those facts as better or worse.¹ But even merely recognizing this heretofore missing

piece of the puzzle transforms our understanding of many central aspects of our cognitive experience.

The stronger version, which I explore but ultimately don't subscribe to, makes a revisionary rather than an additive claim: it proposes that the mind is *at bottom* evaluative in nature. This is to say that the mind's evaluative processes are conceptually *prior* to its perceptual, cognitive, or motor processes. In this sense, the stronger thesis is a type of grand unifying theory for understanding the mind. Notably, the stronger version is related to but distinct from the so-called 'reward is enough' hypothesis, which suggests that reward maximization is sufficient to "drive behavior that exhibits most if not all abilities that are studied in natural and artificial intelligence" (Silver et al., 2021, 1).

Even without the stronger version, reinforcement learning points us to the idea that, as living organisms, we not only continually experience the world, but experience it *as better and worse*. As Haugeland (1979, 619) puts it, the problem with classical computers is that they "don't give a damn." Montague (2006, 19) similarly suggests that the central difference between computers (as we have more traditionally conceived of them) and brains is that the latter use evolved, efficient computations that "care—or more precisely, [that] have a way to care." In my view, these notions of 'giving a damn' or 'caring' are basically right: minds *assess with respect to some goals*, i.e., they 'care' about how things are going with respect to those goals, be they as central as survival or as mundane as getting coffee.

Still, we need a much more systematic way of working out of what this actually means. Moreover, if we do in fact experience the world in this way—that is, *evaluatively*—then this will have important implications for understanding how many of our cognitive capacities function, e.g., why perception and attention select as they do; and, equally, why these capacities break down as they do, e.g., how Major Depressive Disorder may involve *both* a reduction in the primary sensitivity to rewards *and* an individual's reduced ability to *learn* from reward (Huys et al., 2013). Developing this picture is the work I aim to do here.

I build my argument out over stages. For precision, I make several assumptions about the nature of reinforcement learning and its instantiation in minds like ours. I sketch these assumptions, together with their relationship to other versions of reinforcement learning, in Section 16.2. I then briefly survey some of the empirical evidence suggesting that the reinforcement learning paradigm captures something important about biological minds like ours.

In Section 16.3, I get more specific about what that 'something important' is. I do so by characterizing the nature of valuation in the mind, defending the function of valuation as guiding selection and providing evidence for the ubiquity of valuation as selection across a wide range of 'low-' and 'high-level' human psychological capacities.

In Section 16.4, I defend the weaker version of the evaluative thesis. I sketch what we might expect from a strictly 'thinking' mind on the one hand, and from a thinking,

evaluative mind on the other. I suggest that we find plenty of evidence for the latter in a variety of cognitive capacities.

In Section 16.5, I consider the stronger thesis, mapping out how an argument for it might go. I suggest it is a thesis well worth bearing in mind, particularly as we continue to make advancements in artificial intelligence. Nonetheless, I suggest that we presently lack the necessary evidence to subscribe to it wholesale and raise some challenges for securing it going forward.

In Section 16.6, I briefly conclude by addressing what Haugeland calls the common complaint regarding artificial intelligence. According to Haugeland, the complaint suggests that artificial intelligence “pays scant attention to feelings, emotions, ego, imagination, moods, consciousness” (this volume, p. 33). I show how by adopting an evaluative account, we can not only illuminate core aspects of minds like ours, but equally appeal to powerful, computational frameworks to design many (though not all) of the features Haugeland refers to into artificial agents.

To start, let’s look at the narrow end of the argumentative wedge, namely, with a basic sketch of reinforcement learning.

16.2 Program, concepts, and findings

16.2.1 Overview

We can think of reinforcement learning as a research question, as a research program, and as a set of computational tools. As a research question, sometimes called the ‘learning problem,’ reinforcement learning asks how an agent can optimize its behavior by learning from interactions with its environment. For example, how does a baby plover learn the contours of its environment simply by hopping around in it? Or again, how does a newcomer to London find her way around, just by using a map and a bit of trial-and-error? As a research program, reinforcement learning refers to a branch of computer science, together with associated interdisciplinary approaches, that analyzes formal versions of this question and develops computational solutions to it (Dayan and Abbott, 2001; Glimcher and Fehr, 2013). Finally, reinforcement learning methods are the suites of computational algorithms that aim to solve the aforementioned learning problem (Sutton and Barto, 2018a).

As a research program, the reinforcement learning framework makes certain foundational and technical assumptions, with specific versions of the framework committing to some assumptions while suspending or relaxing others. Here, I sketch what I call the ‘reinforcement learning and decision-making’ (RLDM) framework, drawing on assumptions made in both machine learning and computational neuroscience.² Specifically, in addition to assuming many of the somewhat more basic features of the general framework, this version assumes that reinforcement learning is to some degree meaningfully instantiated in the minds of biological organisms, and takes a particular if minimal view regarding the

problem of specifying where rewards come from in biological systems. Throughout, it will be useful to remember that this is just one variant of the general framework among many—though perhaps one that is particularly philosophically useful.

16.2.2 RLDM

Let's start with the basic ingredients. In a reinforcement learning framework, we have an agent and an environment. The agent is the learner or decision-maker in question, and it selects different actions in its environment, where actions can be understood as “any decisions we want to learn how to make,” including mental actions (Sutton and Barto, 2018a, 50). The environment refers to everything ‘outside’ of the agent, which the agent cannot arbitrarily change but rather with which the agent interacts. (For example, in many cases, even parts of the agent's body are considered to be a part of the environment.) The agent and the environment interact in the sense that the agent is presented with sensory information from the environment, and the agent chooses among different actions within the environment (picking what is called a ‘state-action pair’). The environment is then affected by these actions, and the process is (usually) iterated. Notably, the agent may not be able to observe the complete environment and may have no prior knowledge of the environment's dynamics. In addition, the agent may, but by no means needs to, build a model of the environment in order to choose actions in and learn from it.

A distinguishing feature of the reinforcement learning framework is the role of *reward*. Roughly, in reinforcement learning, the agent's objective in the environment is to maximize the cumulative reward it receives over time, where rewards are passed from the environment to the agent. In their influential text, Sutton and Barto call this framing the *reward hypothesis*, specifying, “all of what we mean by [an agent's] goals and purposes can be well thought of as the maximization of expected value of the cumulative sum of a received scalar signal (called reward)” (Sutton and Barto, 2018a, 53). That is, the agent's objective is to maximize its yield of reward as it acts in the world. This objective is characterized by assigning a quantity of intrinsic desirability to each state (or to taking each action in each state). This intrinsic desirability is known as the reward.³

This intrinsic desirability assigned to each state (or taking each action in each state), or *reward*, can be contrasted with the notion of *value*, which captures the expected, discounted, sum of future reward associated with each state (or each action in each state), conditional on a certain policy of action. We can elucidate the distinction between reward and value further using an example adapted from Silver (2015). Imagine an agent trying to find a door. Upon arriving at the door for the first time, the agent receives a *reward* from the environment. But this reward can also be used to assess how relatively good (valuable) individual states are expected to be to the extent that, conditional on a certain action policy, they *lead to* the door and hence the reward. Hence, an agent's ongoing interactions with its environment enable it to continually revise the value attributed to a given state

or state-action pair conditional on a certain policy, upgrading or downgrading as needed. This enables the agent to *learn* the most appropriate actions in the most appropriate states to maximize cumulative reward over time, conditional on a certain policy, in spite of the fact that states (or state-action pairs) can be of high *value* without being intrinsically worthwhile (i.e., *rewarding*).

We can take the example of making and having coffee to help illustrate the difference between a state of high expected value that is nonetheless not technically rewarding. Although only drinking a cup of coffee itself may be intrinsically worthwhile (rewarding), and the grinding of the beans almost certainly is not, the state-action pair of grinding the coffee is nonetheless associated with *expected value*, as it is, conditional on a certain policy, a necessary step or state-action pair on the way to having the coffee.

The distinction between valuable and rewarding states partly helps explain why not every state in an environment needs to be directly rewarding in order for an agent to act appropriately within it.

As a branch of machine learning, reinforcement learning represents the foregoing conceptual features in computational terms. There are countless reinforcement learning algorithms, each with a distinctive computational profile. For example, the *temporal-difference learning algorithm* (TD) represents a computationally efficient way of making predictions about reward in the future. One way to improve predictions over time is to make a prediction about an actual outcome, compare the difference (or error) between the two, and then update the estimates that led to the initial prediction.

To borrow an example from Sutton (1988, 10), suppose you are a weather forecaster in a monotonous climate, charged with making a prediction each week about the chance of rain on the coming Saturday. Each week, you gain more information about the local weather patterns, allowing you to refine your predictive powers. That information could be used in different ways. You could make a prediction on Monday about the weather on Saturday, wait until Saturday, and then update Monday's prediction based on the difference between Monday's prediction and Saturday's actual weather. The temporal difference approach does something a little neater by updating its predictions throughout the week. Having made a prediction Monday about the weather on Saturday, TD lets you compare Monday's prediction to *Tuesday's* prediction about Saturday and adjust your Monday predictions accordingly. For instance, if Monday's prediction for Saturday is a 90% chance of rain, but Tuesday's prediction for Saturday is only a 60% chance, then the temporal difference approach is to lower the Monday prediction for subsequent weeks with similar indicators.⁴

Given that different problem settings present different challenges, there are myriad different RL algorithms in use today. These trade off factors such as memory consumption, computation cost, data efficiency, and stability; some are useful for very small environments, and others are useful for very large environments; some for discrete action spaces,

and others for continuous ones.⁵ Thus, ‘reinforcement learning’ refers to a general learning problem and a suite of computational algorithms, as well as to the branch of computer science devoted to studying them, rather than to any token solution to the problem.

The RLDM version of reinforcement learning adds two assumptions to the basic reinforcement learning framework sketched above. First, it assumes a relationship between reinforcement learning and the minds of biological creatures like us. This assumption is by no means universally held: researchers in machine learning can pursue decades of research and remain entirely agnostic regarding the role of reinforcement learning in biological agents. Similarly, cognitive and comparative psychologists can study the nature of learning and behavior without any appeals to the reinforcement learning framework. However, RLDM follows computational neuroscientists and other decision scientists who suspect that reinforcement learning does, in fact, capture something special about minds like ours. As Dayan and Niv put it, reinforcement learning appears to offer

More than just a computational, ‘approximate ideal learner’ theory for affective decision-making. [Reinforcement learning] algorithms, such as the temporal difference (TD) learning rule, appear to be directly instantiated in neural mechanisms, such as the phasic activity of dopamine neurons. That [reinforcement learning] appears to be so transparently embedded has made it possible to use it in a much more immediate way to make hypotheses about, and retrodictive and predictive interpretations of, a wealth of behavioral and neural data collected in a huge range of paradigms and systems. Dayan and Niv (2008, 1)

Notably, we are free to relax the condition that reinforcement learning is directly *instantiated* in the workings of the brain. It is sufficient to say that reinforcement learning provides remarkably useful frameworks for thinking about decision-making and selection in the mind.

RLDM’s second assumption has to do with the *subjective* nature of reward. As noted above, in the basic reinforcement learning framework, rewards are passed from the environment to the agent when an agent enters certain states of the environment, or when the agent takes certain actions in certain states. This external nature of reward is unproblematic in the context of machine learning because the reward is simply designed by the researcher as a means of communicating what the researcher wants the artificial agent to achieve. But things get thornier when considering biological organisms, since it’s not clear where rewards come from. This question regarding the origin of reward in biology generates what Juechems and Summerfield (2019) call the *paradox of reward*. The issue is paradoxical, the authors contend, because,

No external entity exists that can directly quantify the consequences of each action, like the points that are awarded in a video game for completing levels or shooting monsters. Nor is it obvious that biological systems have a dedicated channel for receipt of external rewards that is distinct from the classical senses. Rather, rewards and punishments are sensory observations—the taste of an apple, the warmth of an embrace—and so stimulus value

must be inferred by the agent, not conferred by the world. In other words, rewards must be intrinsic, not extrinsic.” (2019, 837-838)

Exactly how this conversion between sensory observations and assignments of intrinsic rewards occurs—assuming that it occurs at all—remains the subject of lively theoretical debate. One possible explanation is that minds like ours have evolved specific mechanisms that convert sensory observations into hedonic signals (e.g., see Schultz, 2015). Another, complementary possibility is that, in addition to the evolved mechanisms for basic rewards (e.g., food and water), human beings develop **cognitive setpoints**, akin to homeostatic setpoints, on which reward amounts to a by-product of computing the distance to self-defined goals (e.g., such as getting married or going to graduate school) (Juechems and Summerfield, 2019). Here, RLDM again takes a minimal approach, and merely assumes *that* minds like ours subpersonally assign subjective rewards to, e.g., sensory observations, albeit indirectly; it remains provisionally agnostic about how this assignment takes place.

16.2.3 Substantiating the first assumption

Let’s explore the first assumption in more depth. In what sense does RLDM provide a distinctive, interpretive lens for cognitive neuroscientific evidence?

As gestured at above, arguably the most significant connection is between RLDM and the reward system in the mammalian brain. In the mid-1990s, theoretical and empirical work showed that the firing of dopamine neurons is accurately approximated by the temporal difference learning algorithm (for narrative accounts of the discovery, see (Montague, 2006; Redish, 2013; Colombo, 2014). That is, dopamine neurons fire when an organism experiences a higher- or lower-than-expected value in association with a given state (Schultz et al., 1997). This discovery provides the foundation for the so-called *reward prediction error hypothesis of dopamine neuron activity*, which holds that “one of the functions of the phasic activity of dopamine-producing neurons in mammals is to deliver an error signal between an old and a new estimate of expected future reward to target areas through the brain” (Sutton and Barto, 2018a, 381).

This seminal finding in turn led to the use of reinforcement learning methods to study the neuroscience of vision (Hayhoe and Ballard, 2005; Hikosaka et al., 2006; Hickey et al., 2010), attention (Della Libera and Chelazzi, 2009; Chelazzi et al., 2014; Anderson and Kim, 2018), memory (Patil et al., 2017; Ergo et al., 2020), prospective memory (Krishnan and Shapiro, 1999; Katai et al., 2003; Kliegel et al., 2005; Walter and Meier, 2014), cognitive control (Savine and Braver, 2010; Chiew and Braver, 2014; Cubillo et al., 2019), and above all, decision-making (Sutton and Barto, 2018a; Dayan and Niv, 2008; Rangel et al., 2008; Dayan, 2011; Glimcher and Fehr, 2013).

For example, a systematic body of evidence now indicates that the reward system guides visual fixation and saccadic eye movement, i.e., what we look at, when, and in what order (Liao and Anderson, 2020). Similarly, reward guides what we do or don’t attend to

more precisely than do either location or salience (Anderson and Kim, 2018). Conversely, deficits and disruptions (e.g., by addictive substances) to the reward system are not only implicated in diseases such as Parkinson's and Tourette's, but also in a range of psychiatric disorders, including depression (Huys et al., 2015) and addiction (Hyman, 2005; Redish et al., 2008; Redish, 2013). Arguably, methods from reinforcement learning thus represent an important and, to date, under-utilized framework for elucidating the nature and mechanisms underlying selection between competing states of affairs across a range of 'low'- as well as 'high-level' kinds of cognitive processing.

When proponents say there's something special about RLDM, then, they tend to point to one or both of the following considerations. First, reinforcement learning algorithms successfully predict and characterize the workings of the reward system; by contrast, other approaches, including predictive processing (see Clark, 17 of this volume), often provide merely retrodictive explanations of known phenomena.

Second, the reward system appears to play an outsized role in a range of cognitive capacities, from sensation through to economic choice. The question is, what's the best way of characterizing this role of reward and value in the mind, from a philosophical point of view?⁶

16.3 Valuation

16.3.1 Overview

In principle, the role of reward can be characterized at multiple levels of explanation and across multiple, co-dependent theoretical domains, including in computational terms, cellular- and systems- neuroscientific terms, cognitive neuroscientific and neuroeconomic terms, and psychological and behavioral terms (Hochstein, 2016). For instance, as discussed above, we can capture the role of reward and value in computational terms using methods from reinforcement learning (for an overview, see Sutton and Barto (2018a) though see also hybrid approaches, such as that put forward in Gershman (2015)). Or again, following the Schultz et al. (1997) discovery, we can characterize reward and value in cellular-and-system neuroscientific terms, both in terms of dopaminergic functioning as well as in terms of the more general, system-level neural analyses of the reward system in the brain. At a 'higher' level still, we can characterize reward and value in cognitive neuroscientific and neuroeconomic terms, drawing on behavioral experiments and fMRI data, and using constructs such as 'decision-making,' 'motivation,' and 'willingness-to-pay.' And so on.

In what follows, I characterize the role of reward and value in the mind at roughly a 'conceptual' level of explanation, i.e., at a coarseness of grain typical in the philosophy of mind. Accordingly, my argument also broadens out at this stage, moving from the specifics of RLDM and associated empirical evidence to a more traditional, philosophical

characterization—namely, to characterize a cognitive process I’ll call *valuation*. This is essential for future work in the philosophy of mind, e.g., to enable us to distinguish and understand the relationship between, say, valuation and the philosophical folk psychological notion of *desire* (for work in this spirit, see Schroeder, 2004; Arpaly and Schroeder, 2014), or again, to enable us to distinguish and understand the relationship between, valuation and the various notions of *affect*, *mood*, and *emotion* (for a philosophical discussion of emotion see, e.g., Scarantino and Sousa, 2021).

In this way, the resulting characterization of valuation in some cases *complements* and in some cases *revises* the traditional conceptual machinery used to describe and understand the mind and minds like ours.

16.3.2 Characterising valuation

Recall from the previous section that in basic reinforcement learning, reward is some quantity assigned to represent the intrinsic desirability to each state (or to taking each action in each state), and which is conveyed to an agent when they reach that state. Further, this intrinsic desirability assigned to each state (or taking each action in each state), or reward, can be contrasted with the notion of value, which captures the expected, discounted, sum of future reward associated with each state (or each action in each state), conditional on a certain policy of action. So, while coffee is intrinsically rewarding for me in the morning, grinding coffee or getting milk is not—but these latter states are nonetheless valuable to the degree that, conditional on my action policy, they lead me to my cup of coffee.

Recall in addition that, according to RLDM, the reward hypothesis captures something special about the mind, namely, the substantial role of the reward system in the mammalian brain, where the reward system is itself implicated in a wide range of ‘low-’ and ‘high-level’ cognitive capacities.

I argue that if both of these claims are right, then we can use RLDM and the corresponding empirical evidence to revise our philosophical understanding of what the mind is doing, how it is going about it, and what this kind of processing is for.

Let’s start with the ‘what.’ Very simply, I argue, the mind engages in *valuation*. Informally, I take this to mean that the mind continually attributes reward and value to a range of sensations, perceptions, actions and so on—essentially forming a kind of evaluative layer over the features of its experience.

In more technical terms, I argue that valuation refers to the subpersonal attribution of goal- and context-dependent subjective reward and value to internal and external stimuli. Valuation is *subpersonal* in the sense that it demarcates a causal rather than an intentional mechanism (Dennett, 1969; Drayson, 2014). This is key: the mind routinely, mechanistically assesses states of affairs as better or worse.⁷ Further, it is *goal- and context-dependent* in the sense that what is rewarding or valuable depends on what the agent is trying to do, and when and where the agent is trying to do it. For example, if my goal is to wake up and

have a productive day, then drinking a cup of coffee first thing in the morning is valuable. But if my goal is to rest and get a good night's sleep, then drinking a cup of coffee late at night is not. It is *subjective* in the sense that what is considered rewarding and/or valuable is agent-relative; while this author finds coffee rewarding, many individuals do not. And the term *stimuli* here is intended as a broad catch-all: reward and value can be attributed to external objects (commodities), states, state-action pairs, and action policies, but also to internal states of affairs, such as experiences, feelings, and moods.

In terms of the 'how,' valuation is realized in a number of complementary ways. One important way is through the retroactive attribution of value to states that lead to reward in subsequent states. Recall the task of walking to a nearby door in the previous section. Upon arriving at the door for the first time and therefore receiving or experiencing the *reward*, there occurs a subpersonal, retroactive attribution of *value* to the antecedent states that then led to the reward. That is, there occurs a subpersonal, retroactive attribution of value to the penultimate state, derived from the reward associated with arriving at the 'ultimate' state, i.e., the door. This retroactive attribution in turn continues to feed backwards, i.e., there occurs the subpersonal, retroactive attribution to the antepenultimate state, and so on. In this way, ongoing interactions continue to revise the value attributed to a given state or state-action pair, upgrading or downgrading as needed. For instance, if the baby plover finds a new trove of bugs, the value of a certain path leading to the beach can increase. But values can also be computed 'on the fly' (Balleine and Dickinson, 1998; Langdon et al., 2018), relative to features of context (Hunter and Daw, 2021), and with respect to imagined or expected future states (Gagne and Dayan, 2022; Russek et al., 2021). For instance, if the newcomer to London is traveling from Green Park to Russell Square and Holborn Station is under renovation, the value of taking the blue line decreases.

Here, the main idea is that the mind continually assesses and reassesses states of affairs as better or worse, constructing and casting, to put things in fairly figurative terms, a kind of evaluative fabric over its states and experiences.

16.3.3 Valuation as selection

But it's the 'what for' of valuation that is of most interest (as these things tend to go).

The *function* of valuation in minds like ours, I argue, is to solve for what I call the *selection problem*, or the problem of selecting between one or more competing alternatives. The selection problem can be described in general terms, insofar as the mind must continually select what to *compute*, what to *sense*, what to *perceive*, what to *attend to*, what to *choose* (as an action in the world), and so on. Technically characterized examples of the selection problem include selecting between *multiple action controllers* (Daw et al., 2005), the problem of *perceptual decision-making* (Gold and Shadlen, 2007), and the problem of *action-based decision-making* (Glimcher, 2011). Crucially, as the span of these examples should illustrate, the selection problem occurs ubiquitously in the mind. It occurs at every

major stage of mental processing, from sensation and computation to action, and at every level of description of mental processing, from the sub-personal to the personal.

A central, underappreciated upshot of the RLDM's experiment in mind design, I argue, is that the mind selects between available computations, sensations, perceptions and so on *conditional on* attributions of reward and value.

To illustrate, consider the unlikely phenomenon of binocular rivalry. Binocular rivalry occurs when one stimulus is shown to one eye at the same time as a different stimulus is shown to the other. The resulting experience is of the two images alternating back and forth; *perceptual dominance* in binocular rivalry refers to one of the two images appearing first, or for a longer period of time during the overall duration of the experience of alternation. Notably, both rewarded stimuli and rewarded percepts result in perceptual dominance; that is, participants are more likely to perceive stimuli and percepts associated with a reward (Balceris et al., 2012; Wilbertz et al., 2014; Marx and Einhäuser, 2015; Haas, 2021). Moreover, a complementary phenomenon occurs for punished percepts: participants experience perceptual dominance for the *non*-punished percept in the pair, suggesting that the reward or punishment is *not* simply additional information taken into consideration by Bayes-like predictive processing, as a predictive processing view might suggest (Wilbertz et al., 2014). In this way, *the selection of the perceptually dominant percept is directly conditional on the attribution of reward and value in the binocular rivalry paradigm, i.e., on valuation*. Participants tend to perceive the most rewarded or valuable stimulus or percept. Hence, when it comes to the cognitive task of selecting 'what to perceive,' valuation plays a driving role.

But valuation doesn't just play a driving role in perception. Rather, when I say that the mind is fundamentally evaluative in nature, I mean that we *sense, perceive, and attend* to the features of our environment conditional on our distributions of reward and value attribution, as when we attend to rewarded rather than salient or location-based percepts (Anderson and Kim, 2018). We *remember, remember to remember* (remember prospectively) conditional on reward (for a useful review, see Walter and Meier, 2014). We *allocate our cognitive resources* (in cognitive control) conditional on our distributions of reward and value attributions, as shown by the expected value of control account of cognitive control (Musslick et al., 2015). And we *decide, choose, and plan* our future actions conditional on our distributions of reward and value attributions, as when prior reward experience determines a participant's willingness-to-pay in everyday economic transactions (Plassmann et al., 2007).

Conversely, when the reward system is impaired, for example, through cell death in the basal ganglia (Parkinson's) or due to allostatic shift (substance addiction), there are direct, corresponding deficits in selection: e.g., in motor tremors, mood disorders, and executive dysfunction in Parkinson's disease, and e.g., in cravings, impaired control, and continued

use in spite of overwhelmingly negative consequences in substance addiction (for extended discussions, see Redish, 2004; Redish et al., 2008). And so on.

To emphasize, I do not argue that selection is *synonymous* with valuation. But selection is *conditional* on valuation: we select or avoid what we learn is better or worse over a life-long course of iteration.⁸ Moreover, valuation is deployed and redeployed across a range of selection problems in the mind, including selection in sensation, perception, attention, and cognition generally.⁹

Hence, where I suggested above that the reward system “influences” or is “implicated in” a range of cognitive processing, I can now be much more specific: *valuation guides selection* across the range of mental processing that occurs in minds like ours.

16.4 The weaker thesis

What, then, of the evaluative thesis—the view that the mind is fundamentally evaluative in nature?

At the outset, I suggested that on the weaker version of the view, the mind encompasses *both* thinking *and* evaluation. That is, according to the weaker thesis, the mind does something like ‘see’ two competing stimuli in binocular rivalry, and ‘perceive’ only one of those stimuli at a time, resulting in the signature perceptual experience of perceptual alternation. In a standard case, we might also say that an individual could go on to draw on this perception to form beliefs, draw inferences, and perform all the other kinds of cognitive tasks that are typically associated with, as Haugeland put it, “thinking, intellect,” (Haugeland, Chapter 2, this volume), or as others put it, “intelligence.”

But, on the weaker version of the evaluative thesis, the mind *also* does something else, without which it would not be the mind it is—namely, it continually assesses things as better or worse, conditional on certain goals and aspects of the environment, in the ways described above, i.e., subpersonally, through various forms of attribution, in a two-place relation, and so on.

In this sense, the weaker thesis doesn’t exactly try to unseat the traditional conception of thinking mind but rather complements it by describing a *fundamental* cognitive process that has heretofore been relatively overlooked.

I defend the weaker thesis on three grounds.

First, evidence bears out the positive features of the view. A survey of mature, textbook neuroscience suggests that the reward system is indeed implicated in *basic biophysical processes* such as eating, drinking, and reproduction; in *basic cognitive processes* such as working memory, executive functioning and time estimation; and, crucially, in *all learned behaviors*, ranging from learning-based sensory processing through planning, strategizing, and second-order preference-formation (for a concise review, see Arias-Carrión et al., 2010; for extended discussions, see Glimcher and Fehr, 2013). Equally, the reward system

is implicated in the kinds of ‘*sophisticated*’ *cognitive processes* that are often of interest to philosophers, including in emotional responding, social preference formation, speech and language processing (see especially Simonyan et al. (2012); and also (Ripollés et al., 2014; McNamara and Durso, 2018)), and generalization.

Second, *predictions* made by the weaker thesis are better supported than predictions made by competing theoretical accounts, e.g., by accounts in the predictive processing space or accounts emphasizing the role of emotions in our cognitive processes. Returning to the example of binocular rivalry offers a good example of the former comparison. The weaker thesis predicts that rewards (and negative rewards, i.e., punishments) should influence perceptual dominance in binocular rivalry; predictive processing accounts make no such prediction, and in fact struggle to explain this type of finding post hoc. But as noted above, reward modulates perceptual dominance in binocular rivalry (Haas, 2021).

An example of the latter type of comparison might involve competing explanations of psychopathy. The weaker thesis proposes that psychopathy is a disorder of valuation, perhaps involving an inability to predict negative outcomes, and/or an inability to update appropriately following negative experiences (e.g., see Oba et al., 2021). By contrast, on an account of psychopathy emphasizing emotions, individuals with psychopathic traits fundamentally suffer from a disorder of empathy, or the ability to respond appropriately to emotional stimuli (Hare, 1998; Soderstrom, 2003; Blair, 2007; Brook and Kosson, 2013; Domes et al., 2013; Blair, 2018). Accordingly, the former but not the latter account predicts that individuals with psychopathic traits will exhibit deficits in basic economic decision-making. Here, some evidence seems to bear out the weaker thesis: controlling for other deficits, psychopaths appear to perform significantly worse on the Iowa Gambling Task (Mahmut et al., 2008), as well as on other types of risky decision-making (Takahashi et al., 2014).

Third, *deficits* in the reward system corroborate the view. Here, standard cases again emerge in the computational and cognitive neuroscientific literature, including regarding the aforementioned Parkinson’s and Tourette’s diseases, as well as diseases such as Major Depressive Disorder and different categories of substance addiction. Take the case of prospective memory, or the ability to ‘remember to remember.’ I suggested above that, like so many of our cognitive capacities, prospective memory is conditional on valuation; we are more likely to ‘remember to remember’ something in the future when it’s associated with a reward. For example, participants show higher prospective memory performance for tasks that were associated with a monetary reward as compared to those that were not (Krishnan and Shapiro, 1999). By extension, consistent with the weaker thesis, we would expect to see *deficits* on prospective memory tasks among individuals with Parkinson’s disease. The reasoning goes like this: prospective memory is conditional on valuation, valuation by realized in the reward system in the brain, and the reward system is compromised in

Parkinson's disease. Hence, we should expect deficits on prospective memory tasks among individuals with Parkinson's.

And this is indeed what we find. Individuals with Parkinson's exhibit impairment in several core stages of prospective memory, most notably when it comes to the phases of intention formation and intention initiation (Katai et al., 2003; Kliegel et al., 2005, 2011; Pirogovsky et al., 2012; Ramanan and Kumar, 2013; D'Iorio et al., 2019; Coundouris et al., 2020; though see Zabberoni et al., 2017; Kinsella et al., 2018). Analogous arguments propose that impaired reward valuation, i.e., the dysfunctional underestimation, downgrading, or failure to update regarding rewards in individual with Major Depressive Disorder (Tamura et al., 2017; Rupprechter et al., 2018, 2021), may explain why this demographic also exhibit systematic deficits in prospective memory tasks (Altgassen et al., 2009; Chen et al., 2013; Li et al., 2013; McFarland and Vasterling, 2018). And so on. The basic structure of this third kind of argument, then, is to identify a cognitive capacity modulated by valuation; identify a disease that either upregulates or downregulates valuation (via the reward system) and then determine whether, as predicted by the weaker thesis, individuals with the relevant disorder also exhibit deficits on the corresponding cognitive capacity.

Each of the three sets of reasons gives inductive support for the weaker thesis, by giving a *confirming* instance of it. Saying 'valuation is ubiquitous in the mind' is akin to saying 'lots and lots of swans are white.' This means that the weaker thesis can be disconfirmed—namely, by uncovering a meaningful number of instances where cognitive selection is clearly not, at least in part, underwritten by valuational processes. But this is in fact precisely why I defend the weaker thesis. The normative principles originating in RLDM, together with evidence from the decision sciences, enable us to make a principled but nonetheless fundamentally *empirical* claim about a certain process in the mind—where this claim already brings with it significant high-level implications for understanding the workings of the mind.

By contrast, these same principles and evidence, to my mind, will struggle to bear out something conceptually stronger, including a universal claim regarding the role of valuation in the mind, which I discuss next.

16.5 The stronger thesis

Whereas the weaker thesis holds that valuation is empirically ubiquitous in the mind, the stronger thesis proposes that the mind is *at bottom* evaluative in nature.

There are a couple of ways of understanding the stronger thesis. It can be understood as the claim that valuation as selection guides *all cognitive selection* in the mind. On this understanding, valuation amounts to grand unifying theory for exploring the nature of the mind. This is the stronger, universal version of the weaker thesis. And it can be understood as the claim that valuation is ontologically prior *to* and thus conceptually necessary for

understanding the mind's perceptual, cognitive, and motor processes. We can call the former claim the *scope commitment* and the latter the *priority commitment*.

Prima facie, one might assume that a proponent of the evaluative thesis and RLDM model of valuation would by extension directly subscribe to one or both of these commitments. As we will see, they may hold some theoretical advantages over the weaker thesis. They are also nominally more in line with the prominent 'Reward is Enough' hypothesis (Silver et al., 2021). Nonetheless, I don't commit to either.

So, why not go whole hog and defend the stronger version of the evaluative view? Let's start with the scope commitment.

16.5.1 The scope commitment

The scope commitment is a supercharged version of the weaker thesis. Whereas the weaker thesis holds that valuation is ubiquitous in the mind, the scope commitment holds that valuation lies at the heart of *all* cognitive capacities. Hence, where the weaker thesis suggests that 'lots of swans are white,' the scope commitment rounds up to claim that 'all swans are white,' period.

So formulated, the central challenge with the scope commitment should quickly become obvious: the scope commitment requires defending a universal claim, and no amount of evidence will get us there, as there's always the possibility of an untested counterexample somewhere.¹⁰ The scope commitment is just too easily falsified.

Moreover, it simply doesn't strike me as *likely* that valuation underwrites everything of interest in the mind. The evolved mind is a messy artefact, and at a bare minimum, we can expect 'spandrel' capacities that don't rely on valuation in any interesting sense. I can get plenty of mileage out of the weaker thesis without needing to extend it to the logical limit.

16.5.2 The priority commitment

This leaves us with the priority commitment. The priority commitment is trickier to deal with. The priority commitment makes an ontological claim about the mind: our 'thinking' processes are *conditional* on our evaluative processes. That is, we have the memories, beliefs and so on that we do *in virtue* of our assessments of better or worse. Note that this is analogous to action-first theories in cognitive science; for a review, see (Briscoe and Grush, 2020).

To take a concrete example of this kind of theorizing, one might argue that the normative function of episodic memory is not to encode a past event 'as it actually happened,' but rather to encode a past event in light of what it might be useful for an agent to remember—and by extension, *do*—in the future.

Adopting the priority commitment enables us to make top down rather than inductive predictions regarding the workings of various cognitive capacities. For instance, to continue

with the case of episodic memory, adopting the scope commitment can help us make predictions about what will and won't be remembered, or why individuals experience flashbulb memories (if indeed they do). On the priority commitment, flashbulb memories may contain such an impressive level of detail because, following a traumatic event, it is not clear which features of the preceding event are most relevant to future action, such that 'all' of them are carried forward for future learning. This interpretation draws a close connection between flashbulb memories and the more general credit assignment problem in reinforcement learning, or the problem of determining which actions lead or led to a given outcome (Minsky, 1961; Sutton and Barto, 2018a).

This kind of hypothesis generation is certainly appealing. It's also pretty tempting to defend the priority of valuation as a way of counteracting the standard emphases placed on computation (and predictive processing!) in the philosophical and cognitive scientific literatures. Still, I stop short of doing so, for two reasons.

First, where the scope commitment is too easily falsified, the priority commitment is, conversely, unfalsifiable. If I can describe any cognitive or behavioral phenomenon of interest in terms of the maximization of reward, it becomes more difficult to test the hypothesis.

Second, 'grand unifying' theories of mind encourage us to recast broad swathes of empirical evidence into a single explanatory framework. However, the resulting explanations are sometimes less than illuminating. Moreover, surely some explanatory richness is lost if everything about the mind is ultimately, say, 'imagination,' 'attention,' or 'prediction-error minimization.' In some cases, these kinds of theories even run the risk of discounting evidence that is at odds with their theoretical commitments (Haas, 2021).

There's no reason to expect that the priority commitment would avoid such a fate. To try and keep to a fine-grained and falsifiable view, I thus stick with the weaker thesis.

16.5.3 'Reward is enough'

Finally, let me draw out a few points of comparison between the evaluative thesis explored in this paper and the prominent and somewhat controversial 'reward is enough' (RIE) hypothesis (Silver et al., 2021). **RIE holds that reward maximization is enough to "drive behavior that exhibits *most if not all* abilities that are studied in natural and artificial intelligence"** (Silver et al., 2021, 1, added emphasis mine). Here, reward is understood in the sense put forward by the basic reinforcement learning framework introduced in Section 16.2.

Like the stronger thesis, RIE involves a couple of different claims. First, RIE makes the *epistemological* claim that reward maximization is *enough to understand* many—if not all—features of intelligence. Implicit in this claim is that reward maximization provides better and richer explanations than other rival scientific theories do. Second, RIE makes the *ontological* claim that intelligent processes *just are* reward maximization processes, where

“intelligence, and its associated abilities, can be understood as subserving the maximization of reward by an agent acting in its environment” (Silver et al., 2021, 5). And third, RIE makes that *causal* claim that reward maximization is sufficient to *drive* the kinds of abilities we associate with behavior, such as gathering nuts or playing Go. According to this last claim, the forms of intelligence “implicitly emerge” through and as a direct result of the process of reward maximization. By extension, the authors contend, “a good reward-maximizing agent, in the service of achieving its goal, could implicitly yield all the abilities associated with intelligence that have been considered in natural and artificial intelligence” (Silver et al., 2021, 5).

What is the relationship between the evaluative thesis and RIE? At least on their face, the stronger thesis’s priority commitment and RIE’s epistemological claim appear consistent: the role of reward provides a unified and valuable way of understanding the mind and the nature of intelligence.

But the evaluative thesis and REI come apart on the ontological and causal fronts. At the end of the day, even the stronger thesis amounts to a pair of claims about the function and scope of a *cognitive process* in the mind. By contrast, RIE suggests that *all intelligence processing* is an expression or *by product* of reward maximization where, at bottom, the pursuit of reward drives the emergence of all other kinds of intelligence. These start to look like two very different kinds of arguments.

This being said, one softening feature of RIE is that it makes a pragmatic bet regarding the role of reward maximization in generating diverse forms of intelligence in artificial agents. That is, the authors of RIE propose that pure reinforcement learning frameworks will be sufficient to arrive at artificial general intelligence, without the need for handcrafting or pre-training. The authors acknowledge,

We do not offer any *theoretical* guarantee on the sample efficiency of the reinforcement learning agent. Indeed, the rate at and degree to which abilities emerge will depend upon the specific environment, learning algorithm, and inductive biases; furthermore one may construct artificial environments in which learning will fail. Instead, we conjecture that the solution strategy of learning to maximize reward via interaction will be ‘enough’ for intelligence, and its associated abilities, to emerge in practice. (Silver et al., 2021, 10)

In this sense, by adopting a kind of maker’s approach (Craver, 2021), RIE is at least indirectly falsifiable through efforts to leverage reward maximization to design artificial intelligence.¹¹

16.6 Conclusion

At the outset of this chapter, I proposed that RLDM is an instance of mind design so successful that we have not quite figured out what to do with it yet. I further argued that, in light of this success, we should move beyond characterizing the mind as exhaustively constituted by “thinking, intellect,” as Haugeland originally put it, and begin to recognize

its fundamentally evaluative nature. At the same time, I've sought to distinguish my view, which some philosophers may take to be remarkably strong, from even stronger views, which are more in line with views held by some in the machine learning and reinforcement learning literatures.

By way of conclusion, I want to briefly address what Haugeland called the common complaint about artificial intelligence, namely, that it cannot or may never achieve the rich interiority of everyday life, including “feelings, emotions, ego, imagination, moods, consciousness—the whole ‘phenomenology’ of an inner life. No matter how smart the machines become, there’s still ‘nobody home’ ” (this volume, p. 33). Haugeland’s characterisation is reminiscent of the traditional dichotomized conception of the mind: namely, of understanding the mind in terms of ‘thinking’ and, well, ‘everything else’—even if the ‘everything else’ includes a lot of the important processes.

The notion of valuation—normatively rich, empirically substantiated—allows us to put pressure on this type of traditional, dichotomized view. At a minimum, it challenges the idea that we can in good scientific conscience continue to group together phenomena as disparate as emotions, consciousness, and ego under the heading of ‘phenomenology.’ As noted above, with a notion of valuation in place, we can, for instance, start to work out the relationship and differences between valuation and the various philosophical theories of emotion, or the role of valuation in driving instances of imagination (Gershman et al., 2017). Moreover, without in any way diminishing the ‘thinking’ or ‘computational’ mind, valuation brings with it new avenues for revising our extant philosophical and psychological cognitive taxonomies (Janssen et al., 2017).

More broadly, the notion of valuation challenges our assumptions regarding which aspects of mind can or cannot be quantified—and thereby understood in properly scientific terms. For example, in their discussion of “intelligence” and “intelligent” processes, Silver and colleagues (2021) largely appeal to features of the conventionally thinking mind such as perception, language, and generalization. But what the foregoing discussion should show is that we can also appeal to the normative principles of RLDM to better decompose and understand those allegedly more ‘qualitative’ aspects of the mind such as valuation—and, by extension, our personal-level capacities such as motivation, cognitive control, choice, and moral cognition.

We should also carry these insights forward into our ongoing efforts at mind design. That is, as we make advancements toward more sophisticated artificial intelligence and, particularly artificial general intelligence, we can enrich our understanding of the kinds of mental capacities that we can and should include in these efforts—and we should move past the idea of designing only ‘thinking’ machines in the traditional sense.

Notes for Chapter 16

1. Of course, many approaches in the philosophy of mind and cognitive science posit what we might call ‘compound states,’ such as desires, that may be similarly evaluative. But it’s consistent with such views that evaluative compound states are outliers—that “other stuff”—and overshadowed by traditional descriptive computations and belief-like states and processes. The weaker thesis makes a stronger claim, in that it posits widespread evaluative processing at a *fundamental* level and, notably, where evaluative processing modulates even belief-like states and processes. Thanks to Murray Shanahan for pressing me on this point.
2. Name adapted from Gęsiarz and Crockett (2015).
3. Thanks to Neil Rabinowitz for this formulation.
4. For a more detailed discussion, see Sutton and Barto (2018a, Chapter 6, and especially Example 6.1.)
5. Thanks to Neil Rabinowitz for this formulation.
6. This section is indebted to Sutton and Barto (2018a) and, especially, to Neil Rabinowitz.
7. This subpersonal process very likely plays a role in our personal-level experiences of ‘value,’ ‘valuing,’ and ‘values,’ e.g., see foregoing discussion of willingness-to-pay. But the focus throughout the remainder of this paper will be on the nature and workings of the subpersonal process.
8. It is worth emphasizing that valuation needn’t be ‘online’ in order to guide selection. On the contrary, as in the foregoing example of retroactive attribution, selection can and often is informed by past reward and value attributions. And this ‘carried over’ feature of valuation as selection in turn has important implications for the nature of self-regulation and control, insofar as it implies that at least in many cases, we do not have direct, intrapsychic control over our motivational states (see Haas, in prep). Thanks to Neil Rabinowitz for pressing me on this point.
9. And has been for millions of years: see, e.g., the role of reinforcement signaling in *Drosophila* (Waddell (2013); see also Haas and Klein (2020)). Though this is beyond the scope of the current paper, valuation appears to be a highly conserved cognitive process.
10. Thanks to Carl Craver for helping me drill down on this point.
11. Thanks to Neil Rabinowitz for pressing me on this point, and to Sean and Legassick and Hado van Hasselt for helpful discussions of the REI thesis.

