Rationalization is Rational

Target article and Author's Reply to Peer Commentary

NOTE:

This is a preprint of a target article ultimately published in Brain and Behavioral Sciences alongside peer commentary and a reply. A copy of the reply is appended to this preprint, but the peer commentary is not (because I did not author it).

The complete full text (target article + commentaries + reply) has been made freely available by the publisher here:

https://www.cambridge.org/core/services/aop-cambridge-core/content/view/2A13B99ED09BD802C0924D3681FEC55B/S0140525X19001730a.pdf/rationalization_is_rational.pdf

Rationalization is rational

Fiery Cushman
Harvard University

Rationalization occurs when a person has performed an action and then concocts the beliefs and desires that would have made it rational. Then, people often adjust their own beliefs and desires to match the concocted ones. While many studies demonstrate rationalization, and a few theories describe its underlying cognitive mechanisms, we have little understanding of its function. Why is the mind designed to construct *post hoc* rationalizations of its behavior, and then to adopt them? This may accomplish an important task: transferring information between the different kinds of processes and representations that influence our behavior. Human decision-making does not rely on a single process; it is influenced by reason, habit, instinct, norms and so on. Several of these influences are not organized according to rational choice (i.e., computing and maximizing expected value). Rationalization extracts implicit information—true beliefs and useful desires—from the influence of these non-rational systems on behavior. This is a useful fiction: Fiction, because it imputes reason to non-rational psychological processes; Useful, because it can improve subsequent reasoning. More generally, rationalization belongs to the broader class of “representational exchange” mechanisms, which transfer information between many different kinds of psychological representations that guide our behavior. Representational exchange enables us to represent any information in the manner best suited to the particular tasks that require it, balancing accuracy, efficiency and flexibility in thought. The theory of representational exchange reveals connections between rationalization and theory of mind, inverse reinforcement learning, thought experiments, and reflective equilibrium.

Keywords: *Rationalization; Cognitive dissonance; Self-perception; Theory of mind; Inverse reinforcement learning; Habitization; Social learning; Reflective equilibrium; Useful fiction; Representational exchange*

This fight is over.

THE MAN standing there. In the silence. Two unconscious cops at his feet. Blood on his pants. What just happened? How did he do this? And there's THE GUN in his hand. And God, it just feels so natural -- checking it -- stripping it down -- holding it -- aiming it -- like this is something he's done a million times before...

This is something he definitely knows how to do.

The Bourne Identity

1. Introduction

Jason Bourne is an extraordinary man—a special project of the Central Intelligence Agency. Like a robot, he has been programmed with a vast store of actions that are potentially useful to a clandestine agent. He is fluent in a dozen languages; his gut tells him who to trust and who to fear; he drives like an Italian cabbie; he is handy with a gun.

But Bourne faces an extraordinary problem. He has lost his memory, identity, goals and plans—in sum, his ability to make sense of the world and his own place in it. This problem, and his solution, drive the plot of *The Bourne Identity*. He must figure out what

to believe and what to value by making sense of his peculiar, programmed abilities. The very actions that Bourne performs mindlessly—checking a gun, stripping it, holding it, aiming it—are the clues he uses to rebuild his mind. Jason Bourne learns what to think by seeing what he does.

And, in this respect, he is perfectly ordinary. Each of us faces the same problem every day, and each of us grasps for the same solution. We are never fully certain of what to believe and what to value. But, by observing the actions we are “programmed” to perform, we can draw useful inferences—educated guesses about how the world is, and what to want from it. Like Bourne’s, ours is a rational project: to reverse-engineer the design principles of our automatic actions. Mercifully, for us, the stakes are usually lower. Perhaps that is why it’s so fun to watch Jason Bourne: His life is ours, just more so.

1.1 Rationalization

Rationalization takes an action that has already been performed and then concocts the beliefs or desires that would have made it rational. It is, therefore, exactly

the opposite of rational action¹ (Figure 1). Rational action begins with beliefs and desires and then deduces the optimal action to perform—the one that maximizes desires, conditioned on beliefs. If you believe that a man threatens your life, if you want to live, and if you think he can only be stopped with a bullet, then it is rational to shoot him. Rationalization turns this process on its head: First you shoot a man, and from this you conclude that he threatened your life.

A. Rational Action

B. Rationalization

Figure 1: The relationship between (a) rational action and (b) rationalization.

Sensibly or not, people rationalize all the time. Among psychologists, it is one of the most exhaustively documented and relentlessly maligned acts in the human repertoire. Classic topics such as cognitive dissonance (Festinger 1962), emotion misattribution (Schachter & Singer, 1962), appraisal theory (Arnold 1960, Lazarus 1982), self-perception (Bem 1967, Nisbett & Wilson 1977) and confabulation (Gazzaniga 1967) all have rationalization at their heart. More peripherally, it supports confirmation bias (Nickerson 1998), system justification (Jost & Banaji 1994), motivated reasoning (Kunda 1990), culpable control (Alicke 2000), hindsight bias (Christensen-Szalanski et al. 1991), immune neglect (Gilbert et al. 1998), and more.

Some cases of rationalization are easy to explain. Perhaps we are merely attempting to explain our own behavior, inferring its obscure causes (Bem, 1967). This occurs when, for instance, you notice your furtive glances and flushed cheeks and exclaim to yourself: “I’m falling in love!” (Schachter & Singer, 1962; Dutton & Aron, 1974). Other times, we are hoping to convincingly recast our behavior in a more favorable light (“I ate that last cookie so that nobody else would feel awkward about it!”; Tedeschi et al 1971; von Hippel & Trivers, 2011; Mercier & Sperber 2011).

But other cases of rationalization are much harder to explain. In these hard cases people don’t just tell a story, they actually make themselves believe it (Brehm 1956, Sharot et al. 2010, Vinckier et al 2019). In one experiment, for instance, participants were tricked into believing they had made a “subliminal” choice between two vacation destinations, such as Thailand and Greece. People duped into thinking they chose Greece actually began to like it more, while people who thought they had chosen Thailand showed the opposite preference change (Sharot et al. 2010). Similarly, people believe that a lottery ticket is more likely to win as soon as they have bought it (Langer 1975), or that a horse is more likely to win her race as soon as they have bet on it (Knox & Inkster 1968). These cases are hard to see as anything but gross errors. You are supposed to choose Thailand because you preferred it, or a bet on horse because of its odds. How, then, could those choices justify increasing your preference for Thailand, or your belief in the horse’s odds? These cases seem stubbornly irrational. Why do we drink our own Kool-Aid?

Current theories of rationalization explain how it works, identifying the underlying psychological mechanisms. For instance, the theory of cognitive dissonance posits that we revise our preferences (for Thailand) and beliefs (in a horse’s odds) because we are motivated to reduce dissonance between thought and action. But these theories mostly fail to explain why the brain would contain such mechanisms in the first place (but see Mercier & Sperber 2011). In other words, why would natural selection favor a “dissonance reduction motive”? Classic accounts of rationalization are vague—even silent—at this “ultimate” (Tinbergen, 1963) or “computational” (Marr, 1982) level of analysis.

To address this challenge, it helps to return to one of most basic insights of psychological research: Our behavior is influenced by many psychological processes that are (1) unconscious, (2) non-rational, and yet (3) biological adaptive. For instance, our

¹ I am indebted to Daniel Wegner and Joshua Greene for this succinct statement of the mirror-image relationship between rationality and rationalization, borrowed from

their lecture slides for the undergraduate course *Social Psychology*.

behavior is influenced by instincts, habits, and conformity to social norms (Figure 2a). Rationalization, then, may be a mechanism for extracting valuable information from these adaptive choices and then allowing it to influence the network of beliefs and desires that support reasoning (Figure 2b).

Figure 2. Human behavior is influenced by (a) multiple adaptive processes. (b) Rationalization is a method of extracting information from non-rational processes and making it available to and useful for future reasoning.

According to this view, rationalization is not merely designed to infer the underlying causes of our behavior for the sake of explanation (Bem 1967). It is not, for instance, designed to discover our unconscious reasons: hidden beliefs and desires. Rather, it constructs new beliefs and desires where none had existed, in order to extract information from the *non-rational* processes that influence our behavior. In other words, just as Jason Bourne has been programmed by the CIA with a host of useful reflexes, we have all been programmed: by natural selection, by habit learning, by social learning, and so forth. Thus, just as Bourne can observe his automatic behaviors and extract useful information, so can we.

A simple example illustrates the basic idea. Suppose that an infant crawls to high point and then pulls back from the edge by instinct (Gibson & Walk 1960). This action does not reflect a “belief” that heights are dangerous, or the “desire” to avoid falling; rather, the infant pulls back from the edge by instinct

alone (Gendler 2008). But, having performed this action, rationalization seeks to learn from it—first concocting beliefs and desires that could have produced it, and then adopting them. For instance, the infant might conclude “heights are dangerous”, or adopt the desire to avoid them. She is not inferring the actual beliefs or desires that guided her action; rather, she is constructing new ones, imputing false mental states to an instinctual action. Nevertheless, her new beliefs and desires are adaptive—precipices are dangerous, and you should avoid them. This is no accident: Our instincts embody the hard-won lessons of natural selection. They are, therefore, a rich source of information for your rational mind.

As adults, of course, we rarely rationalize in situations as simple as this. For one thing, our behavior is usually the product of multiple influences, not a “pure” effect of reflex, habit, reasoning, etc. For another thing, our minds are not blank slates bereft of prior knowledge and preferences. These facts make adult rationalization more complex, but no less adaptive. Insofar as non-rational processes exert *some* influence on our behavior, and insofar as that influence is adaptive, we can extract useful information by adopting the beliefs and desires that would have made our actions rational.

Rationalization, then, is “rational” in two senses. First, it is adaptive. This doesn’t guarantee that it always benefits a person in every particular case; to the contrary, many psychologists have made their careers by brilliantly illustrating the ways it can fail. Like any process that is generally adaptive it will be occasionally be maladaptive. But on average, over time, it pays.

Second, at a more specific level, rationalization approximates “inverse planning” models of mental state inference (Baker, Saxe & Tenenbaum 2008; Ng & Russell 2000). Cast as a form of Bayesian inference, inverse planning is sometimes regarded as a rational cognitive process. Of course, the particular mechanisms we use to rationalize our behavior are unlikely to conform to a normative, rational standard (such as Bayesian inference) in every detail. Nevertheless, the basic structure of rationalization can be understood as an approximation of something rational.

Our first two goals are to review existing theories of rationalization, and to contrast these with the present account. Finally, this essay presents a theory of “representational exchange” situating rationalization in a broader framework. Representational exchange describes the flow of information between distinct control systems (reasoning, habits, instincts and norms) to facilitate efficient, adaptive choice. This clarifies the overarching adaptive rationale that unifies

rationalization with many other forms of representational exchange, and highlights its connections to inverse reinforcement learning, habitization, theory of mind, social learning, thought experiments, and the philosophical pursuit of reflective equilibrium.

2. Existing accounts of rationalization

The psychological literature on rationalization is large and varied. Currently, three basic approaches dominate: (1) Cognitive dissonance and consonance, (2) self-perception, and (3) persuasion and impression management. These theories each differ from the current proposal, but do not necessarily compete with it. For one thing, different kinds of rationalization may occur in different contexts. More importantly, existing theories mostly describe the mechanisms of rationalization, whereas the present theory addresses its function. These levels of explanation are often complementary (Tinbergen, 1963; Marr 1982). The goal, therefore, is not adjudicate between theories, marshaling data for some and against others. Rather, it is to clearly present each theory, and then to consider their points of convergence and divergence. This discussion focuses squarely on the rationalization of *action*, touching only superficially on other kinds, such as rationalizing one's beliefs or attitudes.

2.1 Dissonance and consonance

Rationalization is usually explained by positing a desire for “consonance” between thought and action. The simplest model of this kind, “balance theory”, (Heider, 1958/2013), posits that people want to achieve harmony among their attitudes. Thus, if you like your spouse and your spouse likes guacamole, you will tend to acquire a taste for guacamole (or, in theory, a distaste for your spouse). Unfortunately, this model's generality is also its Achilles' heel: it is easy to come up with compelling counterexamples, and hard for the theory to explain them away. As Festinger (1999) quipped, “I like chicken, chickens like chicken food, and I don't like chicken food”.

Festinger's (1956, 1962) own theory of rationalization, cognitive dissonance, was therefore more specific. It posits that only certain sets of beliefs,

desires and actions can occupy states of consonance or dissonance. Although Festinger did not describe the “principle of rational action” as such, he grasped it intuitively²: People tend to act in a way that maximizes their desires, consistent with their beliefs. Sets of beliefs, desires and actions that fit this specific principle are consonant; sets that do not are dissonant³. Crucially, Festinger also proposed that the state of dissonance is psychologically aversive, motivating people to achieve consonance. If you have already acted (e.g., shooting a man), of course, it is too late to adjust your action. Instead, consonance must be achieved by adjusting your beliefs or desires (e.g., deciding he must have been a threat). This is the essence of rationalization.

Cognitive dissonance is the best-known psychological theory of rationalization; indeed, it is among the best-known psychological theories of anything. Its two main pillars enjoy strong empirical support. First, people change their preferences and beliefs in order to match actions they have already performed. Specifically, they adopt the beliefs and desires that would have made their past action rational. A classic method that produces such effects, the “free choice paradigm” (Brehm 1956), is still widely used (e.g., Sharot et al. 2009, Sharot et al. 2010, Vinckier et al 2019). People are given a choice between two things, such as a toaster and a radio, that they value roughly equally. The act of choosing one of these things causes them to value that thing more, and the other thing less. This occurs from an early age and in non-human primates (Egan et al. 2007), even when people cannot remember what they chose (Lieberman et al. 2001), and even when the experimenter tricks them into believing they chose something they didn't (Sharot et al. 2010). People's justifications for such choices can be detailed and elaborate (Johansson et al. 2005). Second, experiments confirm that dissonance is psychologically aversive: People say so (Elliot & Devine 1994), and it is also revealed by covert measures such as affect misattribution (Zanna & Cooper, 1974; Losch & Cacciopo 1990), psychophysiology (Losch & Cacciopo 1990; Harmon-Jones et al 1996) and functional neuroimaging (van Veen et al 2009). In sum, the theory is well known for a good reason.

² Festinger and subsequent theorists of cognitive dissonance often wrote of actions that “follow from” or do not “follow from” cognitions, and this vague concept appears to have typically captured the principle of rational action—i.e., in practice, actions were judged by these theorists to “follow from” beliefs and desires when those actions maximize desire satisfaction conditioned upon beliefs, relative to alternative actions.

³ In many cases dissonance could arise among mental states, such as beliefs and desires; he writes, for instance, “A person may think Negroes are just as good as whites but would not want any living in his neighborhood”. Festinger was primarily interested, however, in the relationship between mental states and actions.

Yet, while the theory of cognitive dissonance describes the psychological mechanisms involved in rationalization, it does not offer an ultimate explanation—an answer to the question: “why did it evolve?”. It is akin to a theory that explains why we eat by saying, “we are motivated by hunger.” Although true and important, such a theory is incomplete: It does not explain why the “hunger motive” evolved (i.e., because food provides the raw material for metabolism, and therefore that organisms who possess the hunger motive tend to outcompete those that lack it).

Festinger was merely the first in a long line of theorists to explain rationalization while eliding its adaptive function (Harmon-Jones & Mills, 1999): Aaronson (1968) proposed that dissonance arises most powerfully when actions are incongruent with a person’s *self concept*; Steele (1988; Steele, Spencer and Lynch, 1993) proposed that dissonance is aversive because people feel that apparently irrational actions threaten their *self-image* or *self-worth*; Beggan (1992) extended this concept to the objects people possess, even when not freely chosen. There are still other possibilities: Perhaps dissonance minimizes post-decisional regret⁴. Each of these proposals elaborates on Festinger’s mechanistic account, but they do not offer an ultimate, adaptive explanation for rationalization.

Indeed, among the classic approaches, the descendants of Heider’s balance theory come closest to an adaptive rationale. “Balance” among beliefs and attitudes can be formalized as a form of logical consistency or constraint satisfaction (Shultz & Lepper, 1999; Thagard, 1989; read & Marcus-Newhall 1993). If a person represents that *Socrates is a man*, and that *all men are mortal*, but that *Socrates is immortal*, something must give. A well-designed system will repair such inconsistencies in a manner that makes its representations more accurate (Ackley et al 1987; see also a review by Gawronsky et al 2018)⁵. This idea of “network repair”, or “cognitive consistency”, can be fruitfully applied to networks of interrelated beliefs or desires (Cushman & Paul, in press).

But how could it be applied to classic cases of rationalization, in which a person revises their beliefs and desires to match their own past *action*? Suppose that you desire cookies and believe them to be in the kitchen, but go to the living room. Clearly one adaptive response to this “imbalance” is to correct your action: To go to the kitchen. This applies the

principle of rational action. But now consider the outcome of rationalization: You either decide you didn’t want cookies in the first place, or else convince yourself that they are actually in the living room. Although this achieves coherence of a kind, it certainly does not improve your desires or beliefs. Rather, it takes take a clear error of reasoning and then multiplies it, infecting thought with a pathology of choice.

Rationalization would indeed be counterproductive if our actions were only produced by sound reasons or outright errors. In this case, dissonance would only arise in cases of error—after all, sound reasoning cannot produce actions that violate principles of rationality. And, if rationalization only applied to errors, then it would propagate those error from action back to our beliefs and desires—a counterproductive result.

Crucially, however, our behavior is influenced by sources other than reason and error. Whether innately, through habit learning, or through cultural learning, our behavior is influenced by processes that are adaptive and yet non-rational. Because these processes share the same “ultimate” purpose of reasoning—fitness maximization—the beliefs and desires that support reasoning can be improved by learning from the behavioral influence of non-rational sources.

2.2 Self perception

Bem’s (1967) theory of self perception provides quite a different explanation for rationalization. It denies the two mechanistic pillars of cognitive dissonance: First, that dissonance is psychologically aversive and, second, that underlying desires or beliefs actually change. Rather, it posits that people only change their *perceived* beliefs and desires.

Bem’s original statement of this theory was heavily influenced by behaviorism in two ways. First, he assumed that we have no direct introspective access to the mechanisms that produce our behaviors. Second, he claimed that we receive strong reinforcement from others when we can explain our behavior in terms of mental states, and consequently we often construct such mentalistic explanations. Desiring mentalistic explanations for our behavior, but lacking introspective access, we resort to our best guess: *post hoc* rationalization.

Although intended as an alternative to cognitive dissonance, self perception theory stands in its own

⁴ Thanks to an anonymous reviewer for this suggestion.

⁵ Although such network repair can be used to improve systems of interrelated beliefs and desires, the relevant processes of “update”, “repair” or “dissonance reduction”

may still be infected by cognitive or attitudinal biases (Gawronski et al 2018).

right as a powerful statement: People are unaware of the causes of their behavior, and often attempt to infer these causes by observing their actions. These basic claims are supported by a wealth of experimental research (Miller 1962/1973; Neisser 1967/2014; Gazzaniga, 1967; Nisbett & Wilson, 1974; Wilson, 1999; Haidt, 2001; Devine, 1988; Greenwald & Banaji, 1995). Bem's (1969) statement of this theory may have been influenced by seminal work on emotion misattribution (Schacter & Singer, 1962) and confabulation in split brain patients (Gazzaniga, 1967). Nisbett and Wilson (1974) later condensed these varied insights into three core claims: People are often unaware of the causes of their behavior; self-reports of the causes of behavior are generated by folk-causal theories; therefore, when they correctly report the cause of their behavior, it is usually the result of inference, not introspection.

In sum, there is undeniably something right about the theory of self-perception. Yet, when aimed at the topic of rationalization, it misses two key empirical marks. First, dissonance induces an aversive psychological state that motivates rationalization. Second—and of greatest importance—rationalization changes people's *actual* beliefs and desires, not just their self-perception. For instance, in the free choice paradigm, objects are not just reported to have higher value after being chosen (or lower after rejection), they are actually chosen more often in the future (or less often, after rejection).

Like Festinger, Bem did not address what “ultimate”, adaptive purpose might be fulfilled by cognitive dissonance. Instead, he focused on its “proximate” psychological motivation: The reinforcement of social partners who demand mentalistic explanations of our behaviors. Neither did he squarely address the issue of whether self-perception is usually accurate or inaccurate. In the grip of behaviorism, Bem may have regarded this as beside the point—he was likely skeptical that our behavior relies on structured mental representations at all. Rather, the essence of Bem's claim was simply that we attribute mental states to ourselves by the same processes that we attribute mental states to others.

The present account builds on important insights of self-perception theories (Bem's, and those that followed it), but with two crucial modifications. First, it posits that the function of self “perception” is not merely to satisfy our own curiosity, or that of our peers. It also constructs new beliefs and desires based on information implicit in other adaptive control mechanisms. Thus, “perception” is a misnomer: Rationalization is designed not to accurately infer unconscious mental states, but to construct new ones; it is not a discovery, but a fiction. Second, we can, should and do actually adjust our beliefs and desires to

match this fiction. This is adaptive because reasoning and non-rational processes are ultimately trying to maximize the same goal: biological fitness. In other words, rationalization is a fiction, but a decidedly *useful* one. Mixed right, it can be nourishing to drink your own Kool Aid.

2.3 Responsibility avoidance and impression management

Finally, some theories posit that rationalization is designed not so much to inform others as to *persuade* them, casting your behavior (or other information) in favorable and possibly deceptive light (Tedeschi et al 1971; von Hippel & Trivers, 2011). When risking blame, for instance, we may profess benign motives or faultless naiveté, even at the expense of the truth.

Mercier and Sperber (2011) go so far as to claim that reasoning itself is principally adapted to the problem of changing others' minds, and therefore interpret rationalization as an adaptive solution to the problem of winning arguments (see also Haidt, 2001; Tetlock 2002). Indeed, on their view, reasoning itself is mostly an instance of rationalization. Its goal is to present information to another person in a manner that compels them, by logic or intuition, to accept your conclusion (see also Haidt, 2001).

At first blush, such theories seem to explain only why we *express* rationalization to others, but not why we adjust our own beliefs or desires (Tedeschi et al 1971). Yet, it is also plausible that “true believers” are better deceivers. In other words, the best way to convince others that you shot an (innocent) man for good reason might be to first convince yourself of his guilt (Trivers, 2000; von Hippel & Trivers, 2011).

This family of theories likely explains a part of the function of rationalization. If the present account also explains part, then it is a complementary but largely independent explanation.

3. Rationalization as construction

More than a century of research shows that our behavior is influenced by multiple processes (Thorndike 1898; Kahneman 2011; Dolan & Dayan 2013). One influence is rational planning: considering the likely outcomes of our behavior according to our beliefs, and then choosing the behavior most likely to maximize our desires. Other influences on our behavior, however, are not organized according to the principle of rational action.

A potential function of rationalization, then, is to construct beliefs and desires that are consistent with the adaptive behaviors generated by non-rational processes, and then to adopt them. In other words, like Jason Bourne, rationalization generates new useful

insights by observing the actions we perform “thoughtlessly”. Later we will view this through a Bayesian lens, as an inversion of a generative model of rational action (Baker, Saxe & Tenenbaum 2008) and thus a variety of “inverse reinforcement learning” (Ng & Russell 2000).

In order to explore the logic of rationalization in more detail it helps to focus on three potentially “non-rational” influences on behavior: instincts (innate influences on behavior), conformity to social norms (socially learned influences on behavior) and habits (reinforced behaviors). Although highly simplified, this taxonomy reveals some important insights about the general structure of rationalization, as well as its specific application in different settings. All three discussions depend on the common assumption that non-rational influences on our behavior are nevertheless adaptive. This is a natural assumption, given that our instincts, habits and norms are all processes shaped by adaptive forces: Biological evolution, reinforcement learning, and cultural evolution, respectively.

3.1 Rational action as planning

Before considering how instincts, norms and habits can improve reasoning, we must have a clearer image of how reasoning itself works. Reasoning, sometimes called planning, chooses actions by expected value maximization (Figure 1a). A simplified model of planning has three parts. First, there is a mechanism for learning a causal model of the world, one’s “beliefs”. This model allows you to predict what is likely to occur in different situations, depending in part on your own actions. Second, there is a mechanism that assigns intrinsic value to certain outcomes, one’s “desires” (also sometimes described as “reinforcement” or “reward”). Third, there is a mechanism that chooses actions by maximizing the satisfaction of your desires, given your beliefs. Our next goal is to understand how a system designed this way could extract useful information from instincts, norms and habits.

3.2 Instincts

Instincts are innate influences on behavior, designed by natural selection, that bias certain actions to be performed in the presence of certain stimuli⁶. For instance, humans instinctively drink when they are thirsty; flee from threats or fight them; reject likely pathogens; fall in love with other humans, and so on. Many of these examples involve very abstract actions (“flee”) or stimuli (“threat”). Instincts need not be “low-level” or concrete, or grounded in a single, well-defined neural mechanism. Rather, “instinct” often describes a very abstract kind of innate mental organization. Its key property is just that the relationship between the stimuli and the actions is innate and direct. For instance, the perception of a threat may directly bias action towards flight. This is what makes instincts different from planning, which would instead require a computation like: “fleeing avoids threats, threats might harm me, and I don’t like being harmed”.

Because instincts are shaped by natural selection they tend to increase our biological fitness. Similarly, rational planning is designed to increase biological fitness. This is an important part of why rationalization makes sense: It extracts information from one adaptive system (instinct) and makes it available to another (rational planning). If a person’s instinct is incongruent with her beliefs and desires, adjusting those beliefs and desires to “match” her action may ultimately improve them.

For instance, suppose that a person instinctively recoils from snakes. This instinct is adaptive because many snakes are venomous, but she happens to be unaware of this. Rationalizing her instinct (i.e., attempting to explain her act of recoiling in terms of beliefs and desires), she might adopt the belief that snakes are dangerous. Similarly, she could rationalize her behavior by adopting a general desire to be far from snakes, and this is a useful desire. In either case the outcomes of her future reasoning are improved.

3.3 Rationalization as a form of rational inference

Even if rationalization could possibly generate true beliefs and useful desires, what guarantees that it would do so typically? The specifics of the snake case

⁶ According to this restrictive definition not all innate mental structure is an instinct. For instance, we have innate mechanisms of perception, memory, attention, etc., but these are not instincts because they are not usefully construed as innate mappings from stimuli to action. Similarly, our innate capacity for rational planning is innate, but it is not an “instinct”. Our innate desires are not instincts either

(e.g., the desire not to be thirsty, not to be attacked, not to be sick, and to be loved). These desires bias action, but not directly. Instead, their influence is mediated by reward learning mechanisms like planning and habit. Thus, although we both innately desire not to be thirsty, and also instinctively drink when thirsty, these are distinct and redundant mechanisms.

are suspiciously convenient; this naïf might have concluded instead that snakes breathe fire, shoot crossbows, or dredge up hurtful memories of adolescence. Such beliefs would explain one's instinctive recoiling, but they are false. Or, she might have adopted the desire to avoid all animals, or all things that are long and straight. What processes could ensure that the rationalized beliefs and desires are, in fact, useful ones?

The construction of new beliefs and desires should presumably be structured as a form of rational inference. In Bayesian terms, the posterior beliefs about snakes ("snakes bite" vs. "snakes shoot crossbows") should be sensitive not just to the likelihood of an action (recoiling) given a percept (snake) and candidate beliefs (e.g., "snakes bite" vs. "snakes shoot crossbows"), but also the prior probability of the candidate beliefs, including their compatibility with other beliefs (e.g., animals can't use crossbows; many long, straight things are perfectly safe).

Indeed, these pieces of information may be integrated in a rational manner, according to Bayes rule. This form of inference has been well characterized in models of "inverse planning" (Baker, Saxe & Tenenbaum 2008; Ng & Russel 2000). This brings into focus an important dimension of the claim that *rationalization is rational*—it is not just biologically adaptive, it may also approximate a well understood form of rational inference. Importantly, however, the "approximation" of a rational inference (at Marr's "computational" level) may be quite cognitively simple (at Marr's "algorithmic" level). These relationships, between "rational inference" and the actual mechanisms of rationalization, are discussed more fully in Section 4.

3.4 Norms

Human psychology is influenced not just by the biological inheritance of natural selection but also by a vast cultural inheritance. And, just as biological natural selection ensures that instincts will typically be adaptive, cultural selection ensures that norms will typically be adaptive (Boyd et al 2011), although maladaptations may arise in each case. Our cultural inheritance takes many forms: concepts (π), artifacts (knives), beliefs (the earth is round), desires (money), norms (drive right, pass left), and much more. The specific form of norms is often transmitted by social conformity. These operate analogously to instincts: just as instincts are innate biases on action, norm conformity may be defined as a set of socially learned biases on action.

As with instincts, norms may be very abstract. For instance, there are cultural norms of cooperation

and fairness that generalize over many diverse features of specific cases. Norms may also be redundant with other kinds of cultural influence. For instance, somebody might comply with the Jewish laws of *kashrut* because (1) she wishes to get along with her religious peers; or (2) she believes that God asks this of her, or (3) it just feels like the right thing to do. These are, in fact, independent and redundant elements of cultural learning. According to our restrictive definition of norms, only the third influence is sufficiently direct to count as a norm. The first two—a desire to get along, and a belief about God—instead influence her behavior indirectly, by reasoning.

Given this homology between instincts and norms, the very same logic that makes instinct a useful target of rationalization therefore also applies to norms. When a person performs a behavior due to cultural influences, she may often be able to extract useful beliefs and desires by rationalizing her action.

Among the indigenous people of Fiji, for instance, it is taboo to eat certain seafoods when pregnant or nursing (Henrich & Henrich, 2010). Most of the taboo seafoods are toxic and pose special risks to fetuses and infants, but the people of Fiji do not have precise knowledge of this—indeed, the taboo extends to several closely related seafoods that are actually harmless. Rather, most mothers avoid eating these foods simply because of the norm. Rationalization, however, might lead a mother to the correct belief that these fish are dangerous. Then, by reasoning, she may generalize other useful conclusions: don't even touch these fish, don't let your baby eat them, don't feed them to sick people, etc.

3.5 A "useful fiction": How theory of mind supports social learning

Because norms are commonly transmitted via observation and imitation (Cialdini & Goldstein 2004; Boyd et al 2011), the benefits of rationalization one's *own* behavior can also be obtained by rationalizing *others'* behaviors. Put more simply, there is a deep homology between rationalization and theory of mind (Figure 3; see also Bem, 1967). This motivates a brief but important detour to consider the relationship between theory of mind, self perception, and rationalization.

Consider again the seafood taboos of Fiji. A mother might conclude that taboo seafood (suppose it is "shellfish") is dangerous to infants by either of two paths: One via rationalization, and another via theory of mind. Following the first path, she first complies with the norm itself, avoiding the shellfish simply because others do. Next, observing her own behavior, she rationalizes that shellfish must be dangerous. Following the second pathway, she instead first

attempts to understand the behavior of his social partners. This act of mental state inference, or theory of mind, really just amounts to rationalization of others' behavior. She concludes that her social partners must believe that shellfish are dangerous. Then, assuming they know something she doesn't, she adopts this belief herself.

Each of these paths involves a crucial step in which a belief is extracted from an action. In the first path it is the observer's own action, so we call it rationalization; in the second path it is another person's action, so we call it theory of mind. (This connection, of course, originates with self perception theory, Bem 1967).

Figure 3: Two parallel pathways of social learning, one characterized by theory of mind and another characterized by rationalization.

This clarifies two important but distinct functions of theory of mind. Many past treatments assume that theory of mind is designed to infer the *true* causes of another person's behavior (Dennett 1987, Baron-Cohen 1995, Gopnik & Meltzoff 1997). Likewise, self-perception theory might be construed as an attempt to accurately infer the true cause of one's own behavior (although Bem himself was agnostic on this point). But we have emphasized an alternative function of rationalization: to *construct* representations that are implied by behavior (one's own, or another person's). This process of construction need not result in a perfectly accurate representation of the causes of behavior in order to be useful (a point made by Dennett, 1989, in introducing the "intentional stance"⁷). To the contrary,

rationalization can extract useful beliefs and desires from the influence of non-rational systems—systems whose influence on behavior had nothing to do with those beliefs and desires. Rationalization is, in this sense, a "useful fiction". It takes the form of inference, but with a very different function.

This same function can also apply to theory of mind. Theory of mind may often involve useful fictions, in which we ascribe inaccurate causes to others' behavior—goal-directed plans, based on beliefs and desires—even when those behaviors were produced by non-rational processes⁸. Although inaccurate, such ascriptions could still extract useful information for us: true beliefs and adaptive desires.

This perspective has at least one attractive feature: Useful or not, theory of mind seems to involve a great deal of fiction. Despite widespread consensus that human behavior is not exclusively rational, nearly all studies of mental state inference posit a folk theory of *rational* action. Despite more than 40 years of study, there is virtually no research on folk theories of instinct, habit, reflex, etc.—in other words, a theory-of-the-rest-of-our-minds. Moreover, what little research exists suggests that people interpret other's actions as the product of goal-directed reasoning far more than it actually is the cause (Gershman et al 2016). Similarly, experimental demonstrations of automatic behavior are often surprising to lay audiences, while experimental demonstrations of rational behavior are not. On the "useful fiction" model, this is because theory of mind is not only designed to infer the true causes of a person's behavior, but also to extract useful beliefs and desires from their behavior even when it was caused by non-rational processes. It is, therefore, biased to perceive all behavior as rational, even though much behavior is not.

3.6 Habits

Habits are a third major non-rational influence on behavior. Habits are learned stimulus-response mappings, often reinforced by reward and punishment (reviewed in Dolan & Dayan 2013). For instance, a person might habitually flip the lights on when they walk into a room because it is typically useful (i.e., rewarding). Each time that the behavior is performed and rewarded, habit learning strengthens the mapping from stimulus to response. As a result, executing

⁷ Dennett's point was that the intentional stance can be a useful fiction for the purposes of describing, explaining and predicting certain kinds of agents or systems. This is surely true, and the present proposal identifies a further, distinct "use" of such fictions: To

improve one's own reasoning by *adopting* the fictive beliefs and desires.

⁸ As Dennett (1989) points out, we might even rationally adopt this "intentional stance" towards non-agents, such as the process of natural selection.

habitual action requires little cognitive effort. A person does not have to consider desires (“I need light”) and beliefs (“switches cause light”) in order to derive the value of performing an action; rather, the behavior is habitized based on its value in the past. But, for the same reason, habitual control can be inflexible. You might habitually switch the lights on even though you are walking into the room of a sleeping baby and want it to be dark. Rational planning, in contrast, can flexibly adjust to new or unusual circumstances.

Despite these differences, there are two key similarities between habit and reasoning: Both involve learning from direct experience, and both are sensitive to the same rewards. These similarities make it challenging to explain how the rationalization of habitual action could provide new information to a system of reasoning. The challenge has two parts.

How to improve beliefs. Any experience that trains a habit also ought to inform your beliefs, and thus it is not clear why our habits would imply useful beliefs that we would not already represent explicitly. For instance, the experiences that formed your habit of turning on the lights ought to have also taught you that flipping the switch makes the light turn on—the very belief you need to flip on the lights by reasoning. In this case, there is no extra “information” for their system of goal-directed reasoning to extract.

This first challenge has a few simple replies. First, a person may simply have forgotten certain facts, and yet nevertheless have retained an adaptive habit. We have all had the experience of being asked for our opinion on something—a restaurant, a book, a colleague—and being able to recall the valence of our feelings (“I know I liked it”; “something about him gave me the heebie-jeebies”) without being able to recall what was eaten, read, or spoken. Even after every detail of an experience evaporates from memory, the residue of our attitudes may remain. According to contemporary theories, this residue—the “cached” values of objects, events or actions—are the basis of habits (Daw & Dolan 2014). By rationalizing habits, we can reconstruct the details that most likely explain them.

Second, and relatedly, a person may have failed to ever formulate the relevant belief (e.g., because they were not paying attention) and yet still have acquired an adaptive habit. Just as it is familiar to have forgotten why we loved a movie or distrusted a person, it is equally familiar to never have been quite sure in the first place. This is possible because habits and beliefs are learned by distinct and dissociable processes (Forde et al., 2011; 2013; Knowlton et al., 1996).

How to improve “desires”. The second and more profound challenge is to explain how adaptive *desires* could also be extracted by rationalization. This challenge is harder because both systems—habit and reasoning—are assumed to begin with the very same set of basic desires (“rewards” (Sutton & Barto 1998) or “primary reinforcers” (Kelleher & Gollub 1962), as they are often called). What information, then, could the habitual system encode that would not already be encoded by the reasoning system? Consider a person who habitually eats cake. Rationalizing this behavior, she concludes “I like cake”. Although true, isn’t this information redundant? She ought to have discovered that she likes cake back when she took her first bite.

The answer to this challenge depends on a key insight regarding the nature of value-guided learning and decision-making. Often our behavior is organized sequentially, with early “instrumental” actions chosen because they eventually bring us to “intrinsically” rewarding states of affairs (Bellman 1954). To plow, sow and harvest are instrumentally valuable actions, for instance, because they ultimately bring a rewarding feast. A major challenge, then, is to discover or estimate the instrumental value of various actions. This challenge is especially obvious in games like chess: We are attempting to learn the instrumental value of moves (or sequences of moves), which is defined by their probability of ultimately attaining checkmate.

Habit and reason estimate value in different ways. Habit learning involves a “backwards-looking” assignment of value: We wait until checkmate is achieved, and then reinforce the sequence of moves that brought us there (Bayer & Glimcher 2005, Morris et al 2006, Roesch et al. 2007; Glimcher 2011). In contrast, reasoning involves a “forward-looking” assignment of value: We mentally simulate hypothetical future sequences of moves, attempting to divine whether they are likely to achieve our goal (Sutton & Barto 1998; Dolan & Dayan 2013).

A further benefit of rationalizing habits that depends upon the hierarchical nature of human planning (Norman & Shallice 1986; Badre & Nee 2017, Botvinick 2008). For instance, if our goal is to make coffee, we plan by calling to mind a series of subgoals (grind beans, get filter, heat water, etc.), which may themselves contain subgoals (turn on the faucet, turn on the kettle, etc.). The essential properties of these subgoals are that they are instrumentally valuable given the superordinate goal, and also that they support generalization across diverse circumstances. But discovering this form of instrumental value does not come for free—indeed, a major challenge for current theories of hierarchical planning is to explain how we discover the appropriate

ways to carve a task into subgoals (Botvinick & Weinstein 2014, Botvinick 2008, Sutton et al 1999).

A crucial function of rationalizing habits, then, maybe to translate the instrumental value representations of the habitual system into goal (Keramati et al 2016) or subgoal (Cushman & Morris 2015) representations useful to the goal-directed system. For instance, if a tennis player habitually rushes to volley at the net after serving, this likely reflects the instrumental value of serve-and-volley for winning a point. When rationalizing this behavior she may say, “my goal was to gain an advantage over my opponent while he was on his heels, in order to quickly win the point”. If she internalizes this subgoal, what has she gained? Not a change to the value of winning the point (which both systems represented), or a change to the cached value of serve-and-volley (which the habitual system represented), but a novel subgoal representation (“try to serve-and-volley!”) that can improve future planning.

In sum, because values are hard to accurately estimate, and because habit and reason estimate value in different ways, the “desires” implicated by habitual action may improve our ability to maximize reward by reasoning.

3.7 Hybrid control

Although it is convenient to act as if certain actions are wholly under habitual control, others wholly instinctual, and so on, this is a caricature. Even the simplest targets of rationalization studied in the laboratory—the choice of a toaster over a radio, for instance—are not the product of pure instinct, habit or norm compliance. Rather, they involve at least some degree of conscious, deliberative planning (“let’s see, what could I do with a new toaster...?”). More generally, it is disputed whether “systems” of habit, instinct or norm adherence could be cleanly severed from reasoning at all (Kool et al., 2017; Graybiel, 2008; Dayan, 2012).

Yet, while reasoning often contributes to choice, it rarely operates alone. Rather, most behavior is the result of some form of “approximate planning”—an elaborate background of automatic and non-rational processes that construct a restricted and tractable decision space in which limited rational planning can effectively guide behavior (Kool et al 2017, Huys 2015, Cushman & Morris 2015, Dayan 2012, Gigerenzer & Selten 2002, Keramati et al 2016). Instinct guides our minds away from rationally deliberating the possibility of marrying our siblings (Lieberman et al 2007); habit guides our minds away from the possibility of making coffee by putting bread in the toaster (Morris & Cushman 2017); norms guide our minds away from the possibility of catching a ride

to the airport by stealing a car (Phillips & Cushman 2017). Non-rational processes also structure tractable planning by identifying valuable end states (Keramati 2016) or goals (Morris & Cushman 2015). A person may seek revenge instinctually, and yet plot her revenge by reasoning; she may seek cocaine habitually, and yet plan to get cocaine by reasoning; she may seek to divide her resources fairly due to blind norm adherence, but then reason carefully about how the fairest division could be accomplished.

Thus, even when our behavior is jointly determined by the influence of rational and non-rational processes, there is an opportunity for rationalization to extract useful information from the influence of non-rational processes and translate these into a form useful to the rational system.

3.8 Summary: Rationalization is rational

Instincts, norms and habits shape our behavior in adaptive ways, but not by rational planning based on beliefs or desires. Still, these influences are adaptive: We instinctively recoil from precipices because they *are* dangerous; norm-based food taboos reflect *real* toxins, and habitually flipping a light switch is *usually* a good idea. Rationalization, then, is a useful fiction: When we observe our own behavior, we infer the beliefs and desires that would have been most likely to have caused that behavior, as if it had been an exclusive product of reasoning. Then we adopt those beliefs and desires. This is adaptive because, on average, the new beliefs are true and the new desires promote fitness. For the same reason, theory of mind might often entail a useful fiction as well: By assuming that others’ behaviors are rational (when they are merely adaptive), we can extract useful information. Crucially, whether rationalizing our own action or others’, the process of inferring information from behavior is structured as a rational inference (Baker, Saxe & Tenenbaum, 2009).

In sum, rationalization exchanges representations of “*Do this!*” for representations of the type “*Believe this!*” or “*Desire that!*”. This is a kind of representational exchange: it extracts information implicit in the representations of non-rational systems and transforms it into the format useful to the rational system. The final section expands this view of representational exchange, showing how rationalization is one example of a much broader class of cognitive operations that facilitate the flow of information among distinct systems of behavioral control.

4. A theory of representational exchange

Rationalization extracts information from non-rational systems and makes it available to reasoning. It is apparent how this could improve reasoning, but why would it improve the overall welfare of the organism? In the end, what matters to an organism is not to have true beliefs and useful desires, but to perform the right actions. Insofar as our actions are already appropriately guided by non-rational forces (habits, instincts and norms), what extra advantage do we gain by improving beliefs and desires?

Properly addressing this question leads to a theoretical framework that encompasses far more than rationalization. Rationalization is just one variety of *representational exchange*: the process of translating information from one psychological system, or representational format, into another. And representational exchange is useful for the whole organism because it organizes information in useful ways—ones that best meets its demands when the information is required. For instance, some ways of representing information demand little computation but are relatively inflexible, “getting it right” in only a restricted range of cases. Others require greater computational demands but are more flexible, “getting it right” in a wider range of cases. Representational exchange allows an organism to transform representations of one kind into representations of another, making thought more efficient by balancing the demands of computational effort and flexibility. This more general perspective, a theory of representational exchange, unifies rationalization with many other cognitive operations.

4.1 The structure and function of representational exchange

During rationalization information flows from non-rational systems to rational ones. Could information flow in the opposite direction—from reason to other adaptive systems, or among the other systems themselves (Figure 4)? Several examples come to mind. During “habitization”, choices that were effortful (i.e., rationally planned) become automatic (i.e., habitual). During “norm internalization”, actions that we observed others perform shape our intrinsic preferences. Although traditionally these processes are considered unrelated, they are all forms of representational exchange: The sharing of information between distinct mechanisms of behavioral control.

Representational exchange is useful because distinct mechanisms of behavioral control have different ways of representing information and guiding action, each with unique advantages and

disadvantages. For instance, habits enable rapid, computationally frugal decision-making that is occasionally sub-optimal, whereas planning attains greater optimality at the expense of time and effort. Representational exchange allows us to keep thought efficient—i.e., to attain the most important opportunities for flexibility and generalization, subject to the resource constraint of a limited cognitive capacity. In this manner it fosters “resource rational” cognition (Griffiths, Leider & Goodman, 2015), improving the overall welfare of the organism.

Figure 4: (a) Rationalization is one example of the more general process of (b) representational exchange.

Consider a simple example. For most people, computing $26 + 52$ takes moment of thought, while $25 + 25$ comes easily to mind. This reflects two different cognitive organizations. One system encodes a procedure for addition, and requires effort to derive specific sums (e.g., 26 and 52). Another system encodes a “precompiled” set of sums—roughly, a table in which one looks up the entry “ $25 + 25$ ” and retrieves “50”. The first requires computation; the second merely requires retrieval.

Why do we have two such systems, and why are certain sums represented one way and other sums another? On the one hand, knowing the rules of addition is useful because it compresses an infinitely large mapping of inputs to outputs (i.e., arbitrary sets of numbers to their sums) via a compact rule. This requires far less memory than, for instance, storing a table of precomputed sums. Although effort is required to compute each sum, this is a worthwhile tradeoff as compared with the storage demands of the

tabular representation and the learning demands of acquiring it.

On the other hand, certain sums must be computed far more often than others. If you are a cashier who makes change every day, then you do store at least a small table of common sums: “nickel + 2 dimes = quarter”, “4 quarters = dollar”, etc. These sums are required so frequently that it would be inefficient to compute them anew each time. Instead, it is worth storing a small “cache” of common sums for ready and quick retrieval.

Any resource rational cognitive system must find efficient ways to represent information, managing the competing demands of computational effort, memory, accuracy and flexibility (Griffiths, Leider & Goodman, 2015; Batchelder & Alexander, 2012). We must choose whether to represent procedures or merely their outputs (Sutton 1991, Gershman et al 2014). We must choose when to represent specifics and when, instead, to fall back on generalities (O'Donnell 2015). We must choose when to be exact and when to be satisfied with an approximation (Gigerenzer & Selten 2002, Kahneman 2011, Kool et al 2017, Daw et al. 2005).

It must be rare that we have attained the optimal balance at any given time, and just as rare that the optimal balance could ever be permanent. Rather, as we learn and change—and as our circumstances and the world around us changes—there is a continual demand to adjust the format of the representations that guide our action. This requires mechanisms for representational exchange.

Some forms of representational exchange will “pack” information into compressed forms, storing outputs, abstractions and heuristics in place of procedures, specifics and computations. Other will perform the reverse operation, “unpacking” information by inferring the more detailed and precise information implicit in outputs, abstractions and rules. Viewed from this perspective, rationalization is a particular kind of “unpacking” that occurs in the specific context of choice behavior: specifically, it unpacks “behavior” into beliefs and desires. It belongs to a broader family of cognitive operations that facilitate representational exchange—not just from non-rational systems to rational ones, but among the many systems that contribute to decision-making.

Representational exchange can be situated within a broader taxonomy of operations demanded by a successful “multi-system” cognitive architecture:

1. **Control.** What are the several mechanisms that guide our behavior? What representations and computations do they rely on, and what are the distinctive advantages and disadvantages of each? Much prior research addresses these questions

(reviewed in Dayan & Dolan 2014; Kahneman 2011; Sloman 1996; Squire 2004).

2. **Metacontrol.** Which system, or weighted combination, guides our behavior at any given time? In other words, from moment to moment, how do we decide whether to act habitually, rationally, by instinct, etc? A growing body of contemporary research addresses this question (e.g. Kool et al 2017; Shenhav et al 2013; Daw & Niv 2005; Griffiths et al 2015).
3. **Exchange.** What mechanisms enable the exchange of information between systems (Lombrozo in press; Gershman et al 2013)? For instance, how can a behavior formerly produced by reasoning become habitual, and how can a behavior that was formerly habitual influence subsequent reasoning? This is our present focus.
4. **Exchange control.** How do we decide what, and when, to exchange? Assuming that representational exchange can be beneficial in the long run, but also carries immediate costs, how is cost-benefit analysis performed? And, when systems embody conflicting information, which system gets prioritized? These are important issues for further development, but they are not pursued here.

The next few sections present an account of representational exchange somewhat more formally, drawing connections to current computational models of decision-making used in psychology, neuroscience and computer science. Although the main focus is on representational exchange among decision-making systems, it is clear that the concept applies beyond the domain of decision-making, and some examples are noted at the end.

4.2 The purpose of representational exchange

The purpose of representational exchange is to make an organism more biologically fit by making its decision-making more efficient. Efficiency is a balance of accuracy and effort. Thus, sometimes we increase efficiency by making more “accurate” decisions; other times by making decisions faster or with fewer cognitive resources. Efficiency can be optimized by sharing information across decision-making systems in order to give an individual an array of options: More controlled and accurate thought, or more rapid and automatic thought, depending on the circumstances.

In order to describe representational exchange in more detail it is useful to use some formal concepts and notations. These should highlight useful themes for those who are already familiar with them, but without frustrating those who are not. The purpose is not to offer a formal model of representational exchange, which is well beyond the scope of this article. Rather, it is to establish points of contact with formal models of control and metacontrol developed elsewhere.

We envision an organism's life as a kind of Markov Decision Process. This means that the organism experiences certain *states* of the world, and in each of these states she performs some *actions*. These actions help to determine the next states she experiences, all of which influence her biological fitness. Any individual's mind can thus be characterized by the probabilistic mapping from states to actions, or *policy*. Colloquially, a policy says: "Here is the thing to do in any given situation" (or "the several things you might do, and their associated probabilities"). From the standpoint of natural selection, there is some "optimal policy" that maximizes expected biological fitness. Nobody actually has an optimal policy, but it is a useful ideal to consider: The total set of instructions for life that maximize your chances of biologically fit children. The closer an organism gets to this ideal, the more fit it is.

Figure 5. An idealized model of behavioral control.

As a simplifying assumption, suppose that instinct, norm compliance, habit and planning each dictate their own specific policy to an organism. In other words, instincts would provide you with one set of instructions; habits with another set of instructions, and so on. These are different mechanisms of behavioral "control". In a perfect world every one of these policies would be identical; specifically, they would all encode the optimal policy. In reality, however, different systems are likely to do better or worse in different cases—i.e., to recommend more or less fitness-maximizing actions in different states.

"Metacontrol" is the problem of deciding how to allocate control to one policy or another in any given situation, or how to blend them.

The goal of representational exchange is to improve the individual policies of each system by transferring information between them. This can improve the overall efficiency of decision-making by allowing optimal-but-effortful thought when appropriate, and suboptimal-but-easy thought when appropriate.

4.2 Advantages and disadvantages of control mechanisms

Before asking how these different influences on our behavior might exchange information, greater precision is required on two points: the different formats in which information is represented, and the relative advantages and disadvantages of each format. Briefly addressing these issues will put us in a better position to understand how and why information might be exchanged between representational formats.

Instinct. Instincts are innate mappings from states to actions that emerge regularly in typical development. The advantages of instinct are speed and reliability: they only depends on the development of the organism and not on learning or reasoning, which both take time and are contingent upon unreliable experience. If an organism innately possessed a set of instincts comprising the optimal policy it would have no need for learning. In reality, however, instincts will not encode the optimal policy because the world changes too fast for biological natural selection to keep pace. Other systems of behavioral control (norms, habits and reasoning) are useful precisely because they allow an organism to adjust its policy towards fitness maximization more rapidly—on the timescale of a single organism's life, rather than a multi-generational one.

Instrumental learning. Whereas instincts implement an innately encoded policy, instrumental learning instead learns a policy by attempting to maximize innately specified rewards. Because instrumental learning occurs within an individual's lifetime, it can improve the agent's policy faster than natural selection. In order for instrumental learning to improve an organism's fitness, natural selection must assign *reward* to states or actions that reliably increase fitness: things like consuming food, acquiring resources, reproducing, etc. This often occurs by estimating the instrumental *value* of certain actions—that is, their expected long-run rewards. Current theories of instrumental control tend to divide between two basic ways of estimating value.

Habit. Habit learning is often modeled as a method of estimating the value of every action in every state based on its history of reinforcement (Daw & Doya 2006; Sutton & Barto 1998). The major advantage of habit learning is its low computational demand. First, it only bothers to estimate the value of states and actions that it has actually experienced; for many tasks, this means that the vast majority of conceivable states and actions are ignored. Second, it precompiles (or “caches”) the instrumental value of actions at the time they are performed, and then draws upon this cached value representation when making future decisions. (This is akin to “caching” the solution to $25 + 25$).

Planning. Planning, like habit, is a variety of instrumental control (Daw & Dayan 2014). It estimates the value of actions prospectively, according to the magnitude and probability of reward of their likely outcomes. When this involves searching over a large model of the potential outcomes it is computationally demanding. Deriving value estimates from an internal model has the advantage, however, of making planning flexible. It can simulate the outcomes of actions it has never performed, it can update its value estimates based on new information, and it can also update them based on new specifications of reward. This may be useful when the agent is tasked with planning towards a specific goal, for instance because of a hierarchical task decomposition (Botvinick & Weinstein 2014, Botvinick 2008, Sutton et al 1999, Cushman & Morris 2015) or due to social coordination such as joint intentionality (Tomasello, 2014; Ho et al. 2016; Kleiman-Weiner et al. 2016).

Norms and social learning. Norms are influences on behavior that are learned from others⁹. An extensive literature shows when and why social learning is valuable (reviewed in Boyd & Richardson 2008). The basic premise is simple enough: Because other people are designed to maximize fitness, and the things that improve their fitness will often also improve yours, you can improve your own fitness by copying others.

⁹ Thus far we have acted as if the mind contained some distinct, explicit norm-based policy—i.e., as a mapping from a state (“going to a friends house for dinner”) to an action (“bring a bottle of wine”) that reflected social learning alone. In reality, however, socially learned norms likely influence behavior not through some proprietary representation, but by altering the kinds of representations used by other mechanisms of behavioral control. These

4.3 Varieties of representational exchange

Having reviewed the representational format of several different influences on our behavior, and the advantages and disadvantages of each, we can now consider several mechanisms of representational exchange in greater detail.

Rationalization as inverse reinforcement learning. We have already seen that rationalization translates observed actions into useful beliefs and desires. Our next goal is to re-describe this idea both more formally and more abstractly, revealing useful connections to several literatures.

Natural selection and instrumental learning share a common structure: Both are trying to maximize some “objective” (fitness, or reward) by shaping the actions we take in the environments we encounter. An interesting property of rationalization is that it can use the common notion of “objective” to turn one kind of objective (fitness) into another (reward). In order to see this more clearly, we will begin by representing both processes (natural selection and instrumental learning) identically, as a function:

$$f(\underbrace{\text{objective, environment}}_{\text{inputs}}) = \underbrace{\text{policy}}_{\text{output}}.$$

Later it will be useful to consider the differences between natural selection and instrumental learning. First, however, having defined this function, consider what happens if we flip its “direction”, swapping inputs and outputs while preserving their mappings. This is the *inverse* function:

$$f^{-1}(\underbrace{\text{policy}}_{\text{input}}) = \underbrace{\text{objective, environment}}_{\text{outputs}}.$$

This inverse function describes rationalization in very abstract terms: you *input* a policy (i.e., a set of actions in a states), and you *output* information about the environment and the objective that the agent is trying to maximize. In more ordinary terms, the “inverse function” could observe a person’s behavior (“actions”) and, on this basis, draws inferences about

include the basic constituents of instrumental learning: representations of value, reward, and the world (i.e., your “causal model”; see Morris & Cushman 2017). Thus, for instance, the norm of wine-gifting might be represented in terms of the instrumental value of the gift, or in terms of a belief that such behavior is expected by social partners, etc. The unifying theme is not the nature of the representation, but the manner in which it was acquired: by social learning.

how the world is (“beliefs”), and what is valuable (“desires”). In short, it acts like Jason Bourne.

This basic idea has been widely explored in computer science, where it is often called “inverse reinforcement learning” (IRL) or “inverse optimal control” (Ng & Russell 2000). Whereas a typical problem that artificial intelligence is designed to solve is *choosing actions given an objective* (i.e., reinforcement learning, or optimal control), in some settings it is desirable to solve the inverse problem: Inferring an objective from a set of observed actions. This goal often arises in social (or “multiagent”) settings. For instance, if a programmer wishes for a machine to learn by observing humans, or to predict human behavior, then the programmer might design the machine to try infer the set of rewards that best explains the human’s behavior. Once a set of likely rewards has been observed, the AI can copy human performance by maximizing those rewards itself, or it can predict human behavior by computing which actions would be reward-maximizing for the human.

IRL is easier said than done. In practice, it is often accomplished by some approximation of Bayesian inference. To see how this works, note that we can consider reinforcement learning itself as a probabilistic generative model. That is, given some specification of reward and an environment (comprised of many states, actions, and the transition probabilities between them), reinforcement learning algorithms generate a probability distribution actions—the policy:

$$P(\text{action} \mid \text{reward}, \text{environment}, \text{state}).$$

What IRL seeks, however, is the opposite: the probability of different rewards and environments given an observed action in a given state. This can be computed by inverting the generative model according to Bayes’ rule (for brevity, action, state, reward and environment are now represented by their first letters):

$$P(r, e \mid a, s) \propto P(a \mid s, r, e)P(r, e \mid s).$$

The leftmost term states what we want: Inferences about rewards and environments generated by the observation of what a person does, a , in some state, s . This is proportional to two things we have: The principle of rational action, which derives a policy by reinforcement learning (i.e., $P(a \mid s, r, e)$), and a prior over rewards and environments. Thus, we can guess how the world is, and what is valuable, by inferring the beliefs and desires that would render observed actions rational given our theory of mind.

Notably, a variety of the same Bayesian inversion is essential to computational models of mental state

inference (Baker, Saxe & Tenenbaum 2008). This could be formalized in the language of a Markov Decision Process (“states”, “actions”, “rewards” etc.), but it is more natural to formalize it in the ordinary language of folk psychology (“beliefs” and “desires”). The homology between these formalizations is, however, apparent. We begin with a “generative model” that predicts action on the basis of an agent’s beliefs (i.e., its perception of its current state, as well as general beliefs about its environment) and desires (i.e., rewards):

$$P(\text{action} \mid \text{desires}, \text{beliefs}).$$

This is given by the principle of rational action. Often, however, our goal is to infer unknown beliefs and desires by observing actions. In this case we may invert the generative model according to Bayes’ rule (again, variables are represented by their first letters):

$$P(d, b \mid a) \propto P(a \mid d, b) P(d, b).$$

Thus, we derive a guess about somebody’s beliefs and desires given the actions we have seen them perform, $P(d, b \mid a)$, from capacity to predict their actions based on beliefs and desires by the principle of rational action, $P(a \mid d, b)$, and a prior over beliefs and desires $P(d, b)$. These examples illustrate that IRL and theory of mind are, in essence, the same.

What happens if we extend the logic of these computations to a setting where there are multiple forms of behavioral control: not just reasoning, but also instinct, habit, norms etc.? Crucially, the basic machinery of IRL can work even when an observer assumes a different cognitive architecture than the actor is actually employing. For instance, the actor could be operating with an innate and unchanging policy derived from some process of natural selection (i.e., “instincts”), in which case its objective is to maximize fitness; nevertheless, the observer could attempt to infer an objective function in terms of “rewards” stated within the reinforcement learning framework (as well as the structure of their common environment). This could be a useful fiction if the observer is designed as a reinforcement learning agent herself, and if the “fitness” objective of the actor is relevant to her reinforcement learning problem. Thus, we shall now stop referring to fitness and desires in common terms as “objectives” and instead represent the crucial difference between them: One is a property of the world, and another is a mental state. This divide is real, and yet it is bridged by the act of rationalization.

Consider, for instance, an organism rationalizing an instinct. Natural selection has shaped our

“instinctual policy” not according to beliefs and desires, but according to actual facts about the world and biological fitness:

$$f(\underbrace{\text{fitness, environment}}_{\text{properties of world}}) = \text{policy},$$

where natural selection defines the objective in terms of fitness. Yet, during rationalization, the inverse function computed is:

$$f^{-1}(\text{policy}) = \underbrace{\text{desires, beliefs}}_{\text{mental representations}}.$$

Put in plain words, whereas fitness and actual environmental conditions shaped our instincts, rationalization extracts *desires/rewards* (a mental representation of an objective function that roughly correspond to the objective of fitness) and *beliefs* (a mental representation of the environment that roughly correspond to actual environmental conditions). Inferences about the causes of our actions become a bridge that translates properties of the world into mental representations of those properties. This makes sense because, roughly speaking, the ultimate function of belief is to represent true properties of the world, and the ultimate function of desire is to represent the fitness consequences of these properties.

In sum, rationalization approximates a form of rational inference, and thus can be understood as a variety of IRL at Marr’s “computational” level—its function is to extract useful information from observed actions. This does not imply, however, that rationalization always involves Bayesian inference at a mechanistic level. In some cases it may, but in other cases relatively simple cognitive processes, akin to those identified by Heider and Festinger, may approximate the “rational” inferences described above.

As we shall see next, a benefit of construing the present theory of rationalization in terms of these more formal concepts, and at Marr’s computational level, is that it makes apparent the relationship between rationalization other forms of representational exchange.

Habitization: Cached value and cached policy. During rationalization information is extracted in order to improve reasoning; we next consider several ways in which analogous processes can extract information in order to improve habitual action.

According to several theories, habits can be understood as “cached” representations of instrumental value—i.e., the expected value of actions, in terms of long-run reward. Plausibly, then, useful habits can be constructed by extracting

information about value from the actions selected by other systems. Here, again, we envision this as a process of inverse reinforcement learning, but this time the goal is to derive a *value function* that can be cached for habitual action:

$$P(v | a, s) \propto P(a | s, v)P(v | s).$$

Thus, while certain instincts, or norms, may not themselves depend on any representation of value, still we may update our cached value representations (i.e., habits) by “inferring” the values that are consistent with the actions performed.

In contrast to such “value-based” models of habit, however, some alternative theories posit that habits depend upon cached *policy*—direct stimulus/response mappings, with no representation of value. On this view, habits are chunks of policy “stamped in” through mere repetition (Dezfouli & Balleine 2012, 2013, Dezfouli et al 2014). For instance, by repeatedly tying our shoes in a particular manner, a specific sequence of actions is chunked into an easily retrieved bundle. Such “cached policy” representations introduce a new target for representational exchange. When a person’s action is determined by instinct, planning, or even value-based habitual action, this action may be cached directly as a policy weight—its informational content thus “exchanged” into a new format particularly suited to efficient online execution.

Figure 6: The representational hierarchy of instrumental learning and decision-making. Information can be stored at any level: Cached policy; cached value; or beliefs and rewards.

Thus, while value- and policy- based theories of habits have been viewed as competitors, they may

instead be viewed as complementary representations within a unified scheme (Figure 6). The hallmark of an instrumental system (whether habit or planning) is that rewards shape values, and values shape policy. Planning makes maximal demands on online computation—it must derive value representations from basic representations of reward and the environment, and then derive a policy from those values. Value-based habitual control (“value caching”) requires substantially less online computation—it can derive a policy from a cached value representation. Value-free habitual control (“policy caching”) makes the minimal demands—it simply enacts the stored policy. Thus, value- and policy-based theories of habit need not be considered as rivals, but as distinct points on a common spectrum.

Offline planning as representational exchange. Reasoning is used not only to choose current actions, but also to support “offline planning”—simulating actions, anticipating their consequences, and then caching the resulting values for rapid decision-making in the future (Buckner & Carroll 2007; Davidson et al 2009; Gershman et al 2014; Daw et al 2011; Gershman et al 2017). This is sometimes described as the rational system “training” the habitual system. It was introduced to the reinforcement learning literature as the DYNA architecture (Sutton 1991), and it is spontaneously deployed by humans (Gershman et al 2014).

Imagine, for instance, a downhill skier competing in the Olympics. She is given a few opportunities to walk the length of the course, building a mental model of it. She then returns to her hotel room and repeatedly visualizes the process of skiing the course. During this offline simulation she is able to precompile a habitual policy. As a result, when she actually traverses the course at speeds approaching 80 miles per hour she can quickly execute her policy without online planning.

Recently, there has been some interest in interpreting imagination, hypothetical and counterfactual thinking, and even causal judgment as forms of offline planning (e.g. Morris et al., 2018; Gershman et al 2017; Lombrozo in press; Icard et al 2018). The perspective offered here makes apparent the connections between these traditional areas of psychological research and machine learning methods such as DYNA and Monte Carlo tree search (Browne et al 2012), which similarly involve offline model-based evaluation to improve a cached value or policy representation.

Representational exchange during social learning. During social learning information is exchanged between individuals. Social learning need not also

involve representational exchange, but it often does. In fact, across several diverse literatures on social learning, culture, and norms, one of the most prominent themes is that social learners extract many different kinds of information when observing others. Representational exchange occurs when the kind of representations guiding an actor’s behavior are not the kind extracted by an observer.

In comparative and developmental psychology, observational social learning is often organized into two broad types: imitation and emulation (Tomasello et al 1987; Whiten et al 2009). Imitation occurs when a learner directly copies the overt behavior of a social target. It could be thought of as something like “policy caching”: a direct update to the probability of performing certain actions in certain states. In contrast, emulation occurs when a learner infers the goal behind somebody’s behavior and then adopts only the goal. This allows the learner to design her own policy to attain that goal, potentially by different means. Thus, it involves a form of IRL, or theory of mind, in which the behavioral policy of another individual is used to generate and then adopt new representations of reward or value.

Social psychologists have developed related taxonomies. Many authors have noted that the effects of norm learning can be “deeper” or “shallower” (Kelman 1958; Cialdini & Trost 1998). These can be construed as different varieties of representational exchange (Morris & Cushman 2017). For instance, sometimes people follow a norm because they represent it as something that other people do and care about—so, they comply with the norm in order to get along. This is sometimes called compliance, or “normative conformity”. The norm is represented explicitly and can be thought of as a part of a person’s world model; it influence behavior via planning about the likely consequences of compliance versus non-compliance.

Other times people follow a norm because they believe that the behavior of others tells them something true and important about the world—e.g., “If everyone is avoiding the roast beef there must be something wrong with it.” This is often called informational conformity. It corresponds to a variety of inverse reinforcement learning in which inferences about others’ beliefs become the basis for updating your own. It could be thought of as a “deeper” form of norm compliance because it gives a person a reason to comply with a norm even in the absence of an audience.

Finally, sometimes people “internalize” a norm—that is, they come to directly value whatever it is that the norm prescribes. This might be because norm compliance becomes habitual (i.e., its value is cached) or, more deeply still, because they represent the very

acts implied by the norm as intrinsically rewarding. This final possibility is the most permanent and influential because reward representations are less subject to subsequent update than value representations, and because they exert an influence on both goal-directed and habitual control systems.

In summary, imitation, emulation, compliance, informational conformity and internalization all embody different models of how an organism can update its representations in response to the same social observations (Morris & Cushman 2017). From the perspective of representational exchange, a major goal of the organism is to update the specific representations that will render the social information most useful to its future behavior.

Other forms of social learning. Another form of social learning is instruction: Roughly, one person telling something to another person. This involves an exchange of information between individuals, but it may still maintain the representational format from teacher to learner. Thus, for instance, a teacher might convey her beliefs to a learner, who would then update her own beliefs; or, the transfer might occur from value to value, reward to reward, or policy to policy. Like mere imitation, these cases involve the exchange of information between individuals, but not an exchange of representational formats.

Learning by instruction—i.e., people talking to each other—may be a very important setting for rationalization of a different type, however. Mercier and Sperber (2011) propose that we often rationalize our behavior to other people through explicit verbal communication (e.g., argumentation) in order to attempt to influence their beliefs and desires in ways that are useful to us. This illustrates the way in which the current theory of rationalization and other past theories may explain distinct and complementary aspects of the phenomenon.

A final form of social learning is evaluative feedback, in which a teacher provides rewards and punishments to a learner in order to exploit their capacity for instrumental learning to ultimately shape their policy. This interesting and complex form of representational exchange is, however, beyond our present scope (but see Ho et al 2017).

4.4 *Beyond decision-making: Other forms of representational exchange*

Representational exchange is useful not just for decision-making processes, but in many other areas of cognition as well. The examples of computing “25 + 25” versus “26 + 52”, for instance, do not really belong to the same general category of decision-making as, say, planning a trip to the grocery store, tying one’s

shoes, or leaping away from a snake. Still, we have relatively computationally cheap, “precompiled” knowledge common sums (such as $25 + 25$), and relatively more computationally intensive methods of deriving uncommon sums (such as $26 + 52$). And, there are circumstances in which it will be optimal to exchange information between these formats.

Representational exchange has been well explored in at least one domain that isn’t principally about decision making: “thought experiments” and other forms of imaginative learning (Lombrozo *in press*). These are cases in which an individual has some kind of intuitive knowledge of a phenomenon (e.g., the behavior of a physical system) and uses this intuitive knowledge as a basis for improving their explicit theory of that phenomenon (e.g., a new theory of physics, such as gravity or relativity). This may appear to be a form of alchemy, conjuring gold from iron filings. In reality, it is more akin to rationalization. First, a generative process is trained by experience; for instance, the visual system might learn to anticipate the motion of physical bodies under various conditions. Next, an individual inspects the information that the generative process makes *explicit* (e.g., an intuition about how physical objects will interact), and uses that information to draw rational inferences about information that is merely implicit (e.g., the laws of motion). Just as rationalization extracts structured information from precompiled value representations, a thought experiment about physics extracts structured information from precompiled predictive perceptual representations.

The resulting information—a theory of physics, for instance—is far more computationally expensive to use under many conditions. (In other words, it is harder to determine how an apple will fall from a tree by applying Newton’s laws than by relying on precompiled predictive perceptual representations). But this explicit theory is also far more flexible, allowing us to solve problems for which we have no adequate precompiled predictive perceptual representations such as landing an astronaut on the moon.

Although rationalization, theory of mind and thought experiments might ordinarily be considered very distinct phenomena, they share both structural and functional similarities—an “unpacking” of implicit information compressed into a narrow format. Noticing these similarities may help us to develop an abstract framework for understanding how, why and when people engage in representational exchanges of any kind (see Batchelder & Alexander, 2012).

Reflective equilibrium as bidirectional representational exchange. Thought experiments are common in philosophy, and they are sometimes used

in the process of achieving “reflective equilibrium”. This is a canonical case of representational exchange. When seeking reflective equilibrium, one contrasts intuitive judgments about particular cases with principled rules or reasons that govern those cases, and then seeks the minimal modifications to both that bring them into alignment (Daniels 2003). It is commonly used during moral reasoning: An individual seeks to bring their “intuitive” judgments of particular cases into alignment with a more general normative theory, and this involves revision to both the particular judgments and to the theory itself.

Reflective equilibrium is attractive to philosophers who want to take intuition seriously without giving it absolute priority. They emphasize that intuition is the result of adaptive processes, which could include habit learning, natural selection and cultural evolution, among others (e.g. Railton, 2014). This echoes a basic argument offered here: The policy recommendations of non-rational systems can be used to improve reasoning, at least from the standpoint of fitness-maximization.

Yet, skeptics have reasonably countered that even if intuition tends towards adaptive outcomes, it is at best a heuristic approximation (e.g. Greene 2014). Why, then, ought intuition ever be favored over reasoning? If a person has devoted appropriate mental effort to reasoning, then what *superior* policy recommendation could arise from the “heuristic” representations of non-rational systems?

One simple reply is simply that a person’s beliefs always might be wrong. But this is a shallow response; presumably we are often quite confident in our reasoning, and yet it still conflicts with our intuitions. A deeper reply is that rationalizing non-rational adaptive processes can improve intrinsic rewards maximized by reasoning. By itself, reasoning has no method for questioning or improving its own goals; it is, in Hume’s (1739) words, “a slave to the passions”. Even a system of reasoning possessed of perfect causal knowledge and unconstrained by computational resource might be improved by information from intuition. Of course, “improvement” in this case means “more biologically fit”; whether this is the kind of improvement that philosophers want is a question for philosophers.

This highlights the relationship between rationalization and the naturalistic fallacy (that is, deriving what “ought to be” from what “is”). In a specific setting—adopting new desires for instrumental learning from observations of one’s own or others’ actions—we have considered an adaptive rationale for this inference. Of course, it is not always the case that people do the best (i.e. most fit) thing; on average, however, there is something important to be learned about what one “ought” to do simply by

observing the actions that people, including yourself, actually perform.

4.5 Summary

More than a century of psychological research shows that our behavior is a product of multiple systems. Given that these each have relative advantages and disadvantages, there is good reason to exchange information across representational formats, gradually optimizing the manner in which information is represented. This perspective highlights the common function of many different processes that convert information between diverse representational formats, and across individuals (Figure 7).

Figure 7: Varieties of representational exchange.

A theory of representational exchange suggests several important topics for further development. First, assuming that representational exchange has a computational cost, how is it allocated efficiently? In other words, is there a mechanism of “exchange control” that regulates when and where representational exchanges of different types occur? Second, when information conflicts arise between systems during representational exchange, how are they adjudicated? One obvious prediction is that an individual will not adopt new beliefs or desires when the action produced by a non-rational system can easily be “explained away” as a mistake. For instance, if a person reflexively jumps away from a fake rubber snake, they might not adopt the belief that rubber snakes are dangerous. Rather, they might conclude that this was a “misfiring” of a non-rational system, and discount it as a source of information. It remains to be seen how this inference can be formalized.

5. Conclusion

Why did I do that? We ask ourselves this question often—perhaps more often than we would like. Why do we bother?

Many past approaches suppose that we are motivated by self-discovery, or the desire to explain ourselves to others (e.g., Bem 1967). Possibly, for instance, our behavior is guided by unconscious reasoning. If so, then *Why did I do that?* is equivalent to *What were my reasons?* — we are trying to divine the obscure beliefs and desires that underwrite our unconscious decision-making. Perhaps this is often so.

But human action is also shaped by non-rational forces. In these cases, any answer to the question *Why did I do that?* that invokes belief, desire and reason is at best a useful fiction. Whether or not we realize it, the question we are actually answering is: *What facts would have made that worth doing?* Like an amnesic government agent, we are trying to divine our programmer’s intent—to understand the nature of the world we inhabit and our purpose in it. In these cases, rationalization implements a kind of rational inference. Specifically, we infer an adaptive set of representations that guide subsequent reasoning, based on the behavioral prescriptions of non-rational systems. This inference is valid because reasoning, like non-rational processes, is ultimately designed to maximize biological fitness. It is akin to inverse reinforcement learning, as well as to Bayesian models of theory of mind, and thus it offers a new interpretation of the function of these processes.

Viewed in this light, rationalization is just one example of a broader set of “representational exchange” mechanisms. Our minds are built to exchange information across multiple systems of behavioral control, allowing for an efficient balance of computation and flexibility during decision making. By perceiving the common function of these processes we can better comprehend their structure.

Acknowledgements. Thanks to Adam Bear, Sam Gershman, Dan Gilbert, Joshua Greene, Mark Ho, Max Kleiman-Weiner, Victor Kumar, Michael Littman, Rebecca Saxe, several anonymous reviewers, and members of the Moral Psychology Research Laboratory for their thoughtful, detailed and generous feedback. This work was supported by grant N00014-19-1-2025 from the Office of Naval Research and grant 61061 from the John Templeton Foundation.

5. References

Ackley, D. H., G. E. Hinton, and T. J. Sejnowski (1987). A learning algorithm for boltzmann

- machines. In *Readings in Computer Vision*, pp. 522–533. Elsevier.
- Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin* 126(4), 556.
- Arnold, M. B. (1960). Emotion and Personality.
- Aronson, E. (1968). Dissonance theory: Progress and problems.
- Badre, D. and D. E. Nee (2017). Frontal cortex and the hierarchical control of behavior. *Trends in Cognitive Sciences*.
- Baker, C. L., R. Saxe, and J. B. Tenenbaum (2009). Action understanding as inverse planning. *Cognition*. 113(3), 329–349.
- Baron-Cohen, S. (1997). Mindblindness: An essay on autism and theory of mind. MIT press.
- Batchelder, W. H. and G. E. Alexander (2012). Insight problem solving: A critical examination of the possibility of formal theory. *The Journal of Problem Solving* 5(1), 6.
- Bayer, H. M. and P. W. Glimcher (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron* 47(1), 129–141.
- Beggan, J. K. (1992). On the social nature of nonsocial perception: The mere ownership effect. *Journal of Personality and Social Psychology* 62(2), 229.
- Bellman, R. (1954). The theory of dynamic programming. *Bulletin of the American Mathematical Society* 60(6), 503–515.
- Bem, D. J. (1972). Self-perception theory. In *Advances in Experimental Social Psychology*, Volume 6, pp. 1–62. Elsevier.
- Botvinick, M. and A. Weinstein (2014). Model-based hierarchical reinforcement learning and human action control. *Phil. Trans. R. Soc. B* 369(1655), 20130480.
- Botvinick, M. M. (2008). Hierarchical models of behavior and prefrontal function. *Trends in Cognitive Sciences* 12(5), 201–208.
- Boyd, R., P. J. Richerson, and J. Henrich (2011). The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences* 108, 10918–10925.
- Brehm, J. W. (1956). Post-decision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology* 52(3), 384.
- Browne, C. B., E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton (2012). A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in Games* 4(1), 1–43.
- Buckner, R. L. and D. C. Carroll (2007). Self-projection and the brain. *Trends in Cognitive Sciences* 11(2), 49–57.

- Christensen-Szalanski, J. J. and C. F. Willham (1991). The hindsight bias: A meta-analysis. *Organizational Behavior and Human Decision Processes* 48(1), 147–168.
- Cialdini, R. B. and N. J. Goldstein (2004). Social influence: Compliance and conformity. *Annu. Rev. Psychol.* 55, 591–621.
- Cialdini, R. B. and M. R. Trost (1998). Social influence: Social norms, conformity and compliance.
- Cushman, F. and A. Morris (2015). Habitual control of goal selection in humans. *Proceedings of the National Academy of Sciences* 112(45), 13817–13822.
- Cushman, F. and L. Paul (in press). Are desires interdependent? In *The Oxford Handbook of Moral Psychology*. Oxford University Press.
- Daniels, N. (2003). Reflective equilibrium. *Stanford Encyclopedia of Philosophy*.
- Davidson, T. J., F. Kloosterman, and M. A. Wilson (2009). Hippocampal replay of extended experience. *Neuron* 63(4), 497–507.
- Daw, N. D. (2015). Of goals and habits. *Proceedings of the National Academy of Sciences* 112 (45), 13749–13750.
- Daw, N. D. and P. Dayan (2014). The algorithmic anatomy of model-based evaluation. *Phil. Trans. R. Soc. B* 369(1655), 20130478.
- Daw, N. D. and K. Doya (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology* 16(2), 199–204.
- Daw, N. D., S. J. Gershman, B. Seymour, P. Dayan, and R. J. Dolan (2011). Model-based influences on humans' choices and striatal prediction errors. *Neuron* 69(6), 1204–1215.
- Daw, N. D., Y. Niv, and P. Dayan (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience* 8(12), 1704.
- Dayan, P. (2012). How to set the switches on this thing. *Current Opinion in Neurobiology* 22(6), 1068–1074.
- de Spinoza, B. (1974). *Ethics*. Hafner Press.
- Dennett, D. C. (1978). Three kinds of intentional psychology. *Perspectives in the philosophy of language: A concise anthology*, 163–186.
- Dennett, D. C. (1989). *The intentional stance*. MIT press.
- Deutsch, M. and H. B. Gerard (1955). A study of normative and informational social influences upon individual judgment. *The Journal of Abnormal and Social Psychology* 51(3), 629.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology* 56(1), 5.
- Dezfouli, A. and B. W. Balleine (2012). Habits, action sequences and reinforcement learning. *European Journal of Neuroscience* 35(7), 1036–1051.
- Dezfouli, A. and B. W. Balleine (2013). Actions, action sequences and habits: evidence that goal-directed and habitual action control are hierarchically organized. *PLoS computational Biology* 9(12), e1003364.
- Dezfouli, A., N. W. Lingawi, and B. W. Balleine (2014). Habits as action sequences: hierarchical action control and changes in outcome value. *Phil. Trans. R. Soc. B* 369(1655), 20130482.
- Dolan, R. J. and P. Dayan (2013). Goals and habits in the brain. *Neuron* 80(2), 312–325.
- Dutton, D. G. and A. Aron (1989). Romantic attraction and generalized liking for others who are sources of conflict-based arousal. *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement* 21(3), 246.
- Egan, L. C., P. Bloom, and L. R. Santos (2010). Choice-induced preferences in the absence of choice: Evidence from a blind two choice paradigm with young children and capuchin monkeys. *Journal of Experimental Social Psychology* 46(1), 204–207.
- Elliot, A. J. and P. G. Devine (1994). On the motivational nature of cognitive dissonance: Dissonance as psychological discomfort. *Journal of Personality and Social Psychology* 67(3), 382.
- Festinger, L. (1962). A theory of cognitive dissonance, Volume 2. Stanford university press.
- Festinger, L. (1999). Reflections on cognitive dissonance: 30 years later. *Cognitive dissonance: Progress on a pivotal theory in social psychology*, 381–385.
- Festinger, L., H. Riecken, and S. Schachter (2017). When prophecy fails: A social and psychological study of a modern group that predicted the destruction of the world. Lulu Press, Inc.
- Foerde, K., E. Race, M. Verfaellie, and D. Shohamy (2013). A role for the medial temporal lobe in feedback-driven learning: evidence from amnesia. *Journal of Neuroscience* 33(13), 5698–5704.
- Foerde, K. and D. Shohamy (2011). Feedback timing modulates brain systems for learning in humans. *Journal of Neuroscience* 31(37), 13157–13167.
- Gawronski, B., R. J. Rydell, J. De Houwer, S. M. Brannon, Y. Ye, B. Vervliet, and X. Hu (2018). Contextualized attitude change. In *Advances in Experimental Social Psychology*, Volume 57, pp. 1–52. Elsevier.
- Gazzaniga, M. S. (1967). The split brain in man. *Scientific American* 217(2), 24–29.
- Gershman, S. J., T. Gerstenberg, C. L. Baker, and F. A. Cushman (2016). Plans, habits, and Theory of mind. *PloS one* 11(9), e0162246.

- Gershman, S. J., A. B. Markman, and A. R. Otto (2014). Retrospective revaluation in sequential decision making: A tale of two systems. *Journal of Experimental Psychology: General* 143(1), 182.
- Gershman, S. J., J. Zhou, and C. Kommers (2017). Imaginative reinforcement learning: Computational principles and neural mechanisms. *Journal of Cognitive Neuroscience* 29(12), 2103–2113.
- Gibson, E. J. and R. D. Walk (1960). The "visual cliff". *Scientific American* 202(4), 64–71.
- Gifford, A. (2008).
- Richerson, P. J., & Boyd, R. (2008). *Not by genes alone: How culture transformed human evolution*. University of Chicago press.
- Gendler, T. S. (2008). Alief in action (and reaction). *Mind & Language*, 23(5), 552–585.
- Gigerenzer, G. and R. Selten (2002). *Bounded rationality: The adaptive toolbox*. MIT press.
- Gilbert, D. T., E. C. Pinel, T. D. Wilson, S. J. Blumberg, and T. P. Wheatley (1998). Immune neglect: a source of durability bias in affective forecasting. *Journal of Personality and Social Psychology* 75(3), 617.
- Glimcher, P. W. (2011). Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences* 108, 15647–15654.
- Gopnik, A., A. N. Meltzoff, and P. Bryant (1997). *Words, thoughts, and theories*, Volume 1. MIT Press Cambridge, MA.
- Graybiel, A. M. (2008). Habits, rituals, and the evaluative brain. *Annu. Rev. Neurosci.* 31, 359–387.
- Greene, J. (2014). *Moral tribes: Emotion, reason, and the gap between us and them*. Penguin.
- Greenwald, A. G. and M. R. Banaji (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological Review* 102(1), 4.
- Griffiths, T. L., F. Lieder, and N. D. Goodman (2015). Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science* 7(2), 217–229.
- Haidt, J. (2001). The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological Review* 108(4), 814.
- Harmon-Jones, E., J. Greenberg, S. Solomon, and L. Simon (1996). The effects of mortality salience on intergroup bias between minimal groups. *European Journal of Social Psychology* 26(4), 677–681.
- Harmon-Jones, E. and J. Mills (1999). An introduction to cognitive dissonance theory and an overview of current perspectives on the theory.
- Heider, F. (2013). *The psychology of interpersonal relations*. Psychology Press.
- Henrich, J. and N. Henrich (2010). The evolution of cultural adaptations: Fijian food taboos protect against dangerous marine toxins. *Proceedings of the Royal Society of London B: Biological Sciences* 277(1701), 3715–3724.
- Ho, M. K., J. MacGlashan, A. Greenwald, M. L. Littman, E. M. Hilliard, C. Trimbach, S. Brawner, J. B. Tenenbaum, M. Kleiman-Weiner, and J. L. Austerweil (2016). Feature-based joint planning and norm learning in collaborative games. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.
- Ho, M. K., J. MacGlashan, M. L. Littman, and F. Cushman (2017). Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition* 167, 91–106.
- Hume, D. (1739) *A treatise of human nature*.
- Huys, Q. J., N. Lally, P. Faulkner, N. Eshel, E. Seifritz, S. J. Gershman, P. Dayan, and J. P. Roiser (2015). Interplay of approximate planning strategies. *Proceedings of the National Academy of Sciences* 112(10), 3098–3103.
- Icard, T., Cushman, F., & Knobe, J. (2018). On the instrumental value of hypothetical and counterfactual thought. *Proceedings of the Cognitive Science Society*.
- Johansson, P., L. Hall, S. Sikström, and A. Olsson (2005). Failure to detect mismatches between intention and outcome in a simple decision task. *Science* 310(5745), 116–119.
- Jost, J. T. and M. R. Banaji (1994). The role of stereotyping in system-justification and the production of false consciousness. *British Journal of Social Psychology* 33(1), 1–27.
- Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.
- Kelleher, R. T. and L. R. Gollub (1962). A review of positive conditioned reinforcement. *Journal of the Experimental Analysis of behavior* 5(S4), 543–597.
- Kelley, H. H. (1967). *Attribution theory in social psychology*. In Nebraska symposium on motivation. University of Nebraska Press.
- Kelman, H. C. (1958). Compliance, identification, and internalization three processes of attitude change. *Journal of Conflict Resolution* 2(1), 51–60.
- Keramati, M., P. Smittenaar, R. J. Dolan, and P. Dayan (2016). Adaptive integration of habits into depth-limited planning defines a habitual-goal-directed spectrum. *Proceedings of the National Academy of Sciences* 113(45), 12868–12873.
- Kleiman-Weiner, M., M. K. Ho, J. L. Austerweil, M. L. Littman, and J. B. Tenenbaum (2016). Coordinate to cooperate or compete: abstract goals and joint intentions in social interaction. In *Proceedings of the Cognitive Science Society*.
- Knowlton, B. J., J. A. Mangels, and L. R. Squire (1996). A neostriatal habit learning system in humans. *Science* 273(5280), 1399–1402.

- Knox, R. E. and J. A. Inkster (1968). Postdecision dissonance at post time. *Journal of Personality and Social Psychology* 8(4p1), 319.
- Kool, W., F. A. Cushman, and S. J. Gershman (2017). Competition and cooperation between multiple reinforcement learning systems. In *Goal-directed decision making: Computations and neural circuits*. New York: Elsevier.
- Kool, W., S. J. Gershman, and F. A. Cushman (2017). Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychological Science* 28(9), 1321–1333.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin* 108(3), 480.
- Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology* 32(2), 311.
- Lazarus, R. S. (1982). Thoughts on the relations between emotion and cognition. *American Psychologist* 37(9), 1019.
- Lieberman, D., J. Tooby, and L. Cosmides (2007). The architecture of human kin detection. *Nature* 445(7129), 727.
- Lieberman, M. D., K. N. Ochsner, D. T. Gilbert, and D. L. Schacter (2001). Do amnesics exhibit cognitive dissonance reduction? the role of explicit memory and attention in attitude change. *Psychological Science* 12(2), 135–140.
- Lombrozo, T. (in press). “learning by thinking” in science and everyday life. In *The Scientific Imagination*. Oxford University Press.
- Losch, M. E. and J. T. Cacioppo (1990). Cognitive dissonance may enhance sympathetic tonus, but attitudes are changed to reduce negative affect rather than arousal. *Journal of Experimental Social Psychology* 26(4), 289–304.
- Marr, D. (1982). Vision: A computational investigation into. WH Freeman.
- Mercier, H. and D. Sperber (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences* 34(2), 57–74.
- Miller, G. A. and R. Buckhout (1962/1973). *Psychology: The science of mental life*. Harper & Row.
- Morris, A. and F. Cushman (2017). A common framework for theories of norm compliance. *Social Philosophy and Policy*.
- Morris, A., J. S. Phillips, T. Icard, J. Knobe, T. Gerstenberg, and F. Cushman (2018). Judgments of actual causation approximate the effectiveness of interventions.
- Morris, G., A. Nevet, D. Arkadir, E. Vaadia, and H. Bergman (2006). Midbrain dopamine neurons encode decisions for future action. *Nature Neuroscience* 9(8), 1057.
- Neisser, U. (2014). *Cognitive psychology: Classic edition*. Psychology Press.
- Ng, A. Y., S. J. Russell, et al. (2000). Algorithms for inverse reinforcement learning. In *International Conference on Machine Learning*, pp. 663–670.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology* 2(2), 175.
- Nisbett, R. E. and L. Ross (1980). Human inference: Strategies and shortcomings of social judgment.
- Nisbett, R. E. and T. D. Wilson (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review* 84(3), 231.
- Norman, D. A. and T. Shallice (1986). Attention to action. In *Consciousness and self-regulation*, pp. 1–18. Springer.
- O’Donnell, T. J. (2015). *Productivity and reuse in language: A theory of linguistic computation and storage*. MIT Press.
- Phillips, J. and F. Cushman (2017). Morality constrains the default representation of what is possible. *Proceedings of the National Academy of Sciences* 114(18), 4649–4654.
- Railton, P. (2014). The affective dog and its rational tale: Intuition and attunement. *Ethics* 124 (4), 813–859.
- Rakison, D. H. and J. Derringer (2008). Do infants possess an evolved spider-detection mechanism? *Cognition* 107(1), 381–393.
- Read, S. J. and A. Marcus-Newhall (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology* 65(3), 429.
- Roesch, M. R., D. J. Calu, and G. Schoenbaum (2007). Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. *Nature Neuroscience* 10(12), 1615.
- Saxe, R., S. Carey, and N. Kanwisher (2004). Understanding other minds: linking developmental psychology and functional neuroimaging. *Annu. Rev. Psychol.* 55, 87–124.
- Schachter, S. and J. Singer (1962). Cognitive, social, and physiological determinants of emotional state. *Psychological Review* 69(5), 379.
- Sharot, T., B. De Martino, and R. J. Dolan (2009). How choice reveals and shapes expected hedonic outcome. *Journal of Neuroscience* 29(12), 3760–3765.
- Sharot, T., C. M. Velasquez, and R. J. Dolan (2010). Do decisions shape preference? evidence from blind choice. *Psychological Science* 21(9), 1231–1235.
- Shultz, T. R. and M. R. Lepper (1999). Computer simulation of cognitive dissonance reduction. *Cognitive dissonance: Progress on a pivotal theory in social psychology*, 235–265.
- Singh, S., R. L. Lewis, and A. G. Barto (2009). Where do rewards come from. In *Proceedings of the annual*

- conference of the cognitive science society, pp. 2601–2606.
- Singh, S. P. and R. S. Sutton (1996). Reinforcement learning with replacing eligibility traces. *Machine Learning* 22(1-3), 123–158.
- Steele, C. M. (1988). The psychology of self-affirmation: Sustaining the integrity of the self. In *Advances in Experimental Social Psychology*, Volume 21, pp. 261–302. Elsevier.
- Steele, C. M., S. J. Spencer, and M. Lynch (1993). Self-image resilience and dissonance: The role of affirmational resources. *Journal of Personality and Social Psychology* 64(6), 885.
- Sussman, R. and R. Gifford (2018). Causality in the theory of planned behavior. *Personality and Social Psychology Bulletin*, 0146167218801363.
- Sutton, R. S. (1988). Learning to predict by the methods of temporal differences. *Machine Learning* 3(1), 9–44.
- Sutton, R. S. (1991). Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin* 2(4), 160–163.
- Sutton, R. S. and A. G. Barto (1998). *Reinforcement learning: An introduction*, Volume 1. MIT Press Cambridge.
- Sutton, R. S., D. A. McAllester, S. P. Singh, and Y. Mansour (2000). Policy gradient methods for reinforcement learning with function approximation. In *Advances in neural information processing systems*, pp. 1057–1063.
- Sutton, R. S., D. Precup, and S. Singh (1999). Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence* 112(1-2), 181–211.
- Tedeschi, J. T., B. R. Schlenker, and T. V. Bonoma (1971). Cognitive dissonance: Private rationalization or public spectacle? *American Psychologist* 26(8), 685.
- Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences* 12(3), 435–467.
- Thorndike, E. L. (1898). Animal intelligence: an experimental study of the associative processes in animals. *The Psychological Review* 2(4), i.
- Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für Tierpsychologie* 20(4), 410–433.
- Tomasello, M. (2014). *A natural history of human thinking*. Harvard University Press.
- Tomasello, M., M. Davis-Dasilva, L. Camak, and K. Bard (1987). Observational learning of tool-use by young chimpanzees. *Human Evolution* 2(2), 175–183.
- Trivers, R. (2000). The elements of a scientific theory of self-deception. *Annals of the New York Academy of Sciences* 907(1), 114–131.
- Van Veen, V., M. K. Krug, J. W. Schooler, and C. S. Carter (2009). Neural activity predicts attitude change in cognitive dissonance. *Nature Neuroscience* 12(11), 1469.
- Vinckier, F., L. Rigoux, I. T. Kurniawan, C. Hu, S. Bourgeois-Gironde, J. Daunizeau, and M. Pessiglione (2019). Sour grapes and sweet victories: How actions shape preferences. *PLoS Computational Biology* 15(1), e1006499.
- Von Hippel, W. and R. Trivers (2011). Reflections on self-deception. *Behavioral and Brain Sciences* 34(1), 41–56.
- Watkins, C. J. and P. Dayan (1992). Q-learning. *Machine Learning* 8(3-4), 279–292.
- Whiten, A., N. McGuigan, S. Marshall-Pescini, and L. M. Hopper (2009). Emulation, imitation, over-imitation and the scope of culture for child and chimpanzee. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1528), 2417–2428.
- Wilson, T. D. (2004). *Strangers to ourselves*. Harvard University Press.
- Wilson, T. D. and E. W. Dunn (2004). Self-knowledge: Its limits, value, and potential for improvement. *Annual review of Psychology* 55.
- Zanna, M. P. and J. Cooper (1974). Dissonance and the pill: An attribution approach to studying the arousal properties of dissonance. *Journal of Personality and Social Psychology* 29(5), 703.

Rationalization as representational exchange: Scope and mechanism

Fiery Cushman

Harvard University, Cambridge, MA 02138

Abstract: The commentaries suggest many important improvements to the target article. They clearly distinguish two varieties of rationalization – the traditional “motivated reasoning” model, and the proposed representational exchange model – and show that they have distinct functions and consequences. They describe how representational exchange occurs not only by post hoc rationalization but also by ex ante rationalization and other more dynamic processes. They argue that the social benefits of representational exchange are at least as important as its direct personal benefits. Finally, they construe our search for meaning, purpose, and narrative – both individually and collectively – as a variety of representational exchange. The result is a theory of rationalization as representational exchange both wider in scope and better defined in mechanism.

R1. Introduction

It is a delight to receive such thoughtful commentaries, and a gift to be able to reply. They offer many important improvements to the target article. My goal is to make these as clear as possible. To this end, several sections of this reply summarize the most valuable revisions to the theory of rationalization as representational exchange offered by the commentaries (see R2.1, R4.3, R5.4, and R7.4). The target article focused the representational exchange function of rationalization, but the commentaries show that is important to also acknowledge its self-serving functions and to keep these distinct. The target article focused on post hoc rationalization, but these commentaries extend it to cover antecedent and more dynamic cognitive processes. The target article focused mostly on the personal benefits of representational exchange, but these commentaries extend it to cover the many important interpersonal and collective benefits. Most significantly, several of these commentaries translate the academese of the target article (“rationalization extracts implicit information”) into the native language of human experience. We constantly try to find, or make, meaning in our lives. We do this individually, and we do it collectively. We do it through self-reflection: By interrogating the purpose of our actions and attitudes and by striving for consistency and coherence among our beliefs. And, whether we realize it or not, we do it for self-improvement. The theory of representational exchange offers a natural account of these behaviors. I conclude by considering the challenges, opportunities, and pitfalls of self-improvement by rationalization.

Perhaps, then, I should not have titled the target article “Rationalization Is Rational,” but instead “Self-Reflection Is Useful.” Although I can’t shake the feeling that there was something right about the original title, I’ll have to think a little harder about what that might be.

R2. Motivated reasoning versus representational exchange

There must be at least two kinds of rationalization. One of these is proposed by the target article, and **Ullman** calls it “self-benefitting.” Through representational exchange it improves our beliefs, desires, and ultimately our behavior. But several commentaries urge that we not overlook the other kind of rationalization, which Ullman calls “self-serving” (see also **Quilty-Dunn**), and **Ellis & Schwartzgebel** call “rationalization in the pejorative sense.” It is basically a variety of motivated reasoning. In order to convince ourselves or others that we did the right thing, we invent explanations for our behavior that cast it in favorable light. Such rationalizations do not distill a liquor from the mash of non-rational influences, but instead slap a fraudulent label on a jug of rotgut. Ellis & Schwartzgebel, Quilty-Dunn, **Tierney & Uhlmann**, **Altay & Mercier**, **D’Cruz**, and **Brody & Costa** argue that most rationalization is self-serving. Whether or not it is more common, it is certainly more prototypical. Perhaps Ellis & Schwartzgebel say it best: “Rationalization has a bad name for good reason, and that will be missed on Cushman’s theory.”

How badly is it missed? The target article certainly doesn’t deny that rationalization is sometimes self-serving. To the contrary: “This family of theories likely explains a part of the function of rationalization. If the present account also explains a part, then it is a complementary but largely independent explanation” (target article, sect. 2.3, para. 4). Yet, by calling these two kinds of rationalization by the same name, the target article implied an unnecessary competition between them.

The wiser course, urged by **Ellis & Schwartzgebel**, **Ullman**, **Quilty-Dunn**, and **D’Cruz** is to allow each its own name. Following Ullman, I will call them self-serving (i.e., rationalization as motivated reasoning) and self-benefitting (i.e., rationalization as representational exchange). Self-benefitting rationalization is not necessarily distinct from self-serving rationalization

because of consequences – that is, adaptive versus maladaptive outcomes. Presumably both are usually good for people but sometimes bad. Neither is it necessarily distinct in its mechanisms. It is possible that both forms of rationalization share a common mechanism. Rather, what distinguishes self-benefitting rationalization from self-serving rationalization is its function: representational exchange. Following from this function, self-benefitting rationalization should generally improve the quality of reasoning, while self-serving rationalization will generally diminish it.

Of course, the key question is whether self-benefitting rationalization actually happens. Insofar as the human mind is designed to rationalize, is representational exchange ever the ultimate goal? Several commentaries raise important doubts. For instance, existing research shows that rationalization is often empirically associated with a worse, not better, quality of reasoning and action. Summarizing some this evidence, **Tierney & Uhlmann** entertain the possibility that “rationalizers could tend to be unhappy, socially unpopular underperformers, rejected and ineffective due to their self-serving arguments and lack of insight into their own actions” – a vision more or less shared by **D’Cruz, Brody & Costa**, and Tierney & Uhlmann. (**Quilty-Dunn** shares the view that rationalization degrades reason, but views it as a useful “drive state that pushes us to reorient our beliefs in ways that keep our self-esteem up, even when we are irrational, unstable, or immoral, and thereby keep ourselves motivated by avoiding the awful truth.”)

In support of these claims, **Tierney & Uhlmann** review evidence the rationalization of gender bias in hiring decisions is associated with greater gender bias in hiring decision. **Brody & Costa** review the evidence that rationalization of violence and criminality is associated with greater antisocial behavior. **Quilty-Dunn** reviews evidence that rationalization of meat

consumption is associated with the belief that it is “necessary, normal, natural, and nice.”

If you believe that gender bias, violence, criminality, and eating meat are bad things, then surely it is bad to rationalize them. But does that mean that rationalization itself is bad?

Obviously, good objects can be used for bad ends: A chef’s knife makes a fist fight worse, but it is awfully handy for cooking. When we look for bad effects of rationalization, we can find them; what would happen if we instead looked for good effects? For instance, is rationalizing charitable acts associated with more prosocial behavior?

Brody & Costa review evidence that high levels of rationalization characterize several mental disorders. They conclude that rationalization is generally maladaptive. But it is easy to see the flaw in arguments of this kind. Anxiety disorders are characterized by excessive fear, but this doesn’t mean that fear is generally maladaptive. Presumably many mental disorders arise from the dysregulation of thought processes that, in moderation, are essential to proper psychological function. In fact, as **Saxe & Nettle** note, self-benefitting rationalization (i.e., critical self-reflection about the causes of one’s behavior) is essential to cognitive behavioral therapy, one of the most effective, evidence-based treatments available for mental disorders. It is an advantage of the representational exchange framework that we can identify both why this kind of behavior can be helpful and also how its misapplication can cause harm.

De Neys summarizes further evidence consistent with self-benefitting rationalizations. Across several experiments, they find that rationalization is often used in service of finding the right explanations for the right actions, but in cases where the actions themselves were a product of intuition. For instance, consider the famous “bat and a ball” question that appears in the Cognitive Reflection Task (Frederick 2005). When people report the correct answer to this word problem, they often cannot say exactly why it is correct. Rather, they have to deliberate for a

moment about why it is correct, eventually explaining their answer in terms of sound reasons, but ones that were only rendered explicit by post hoc rationalization. This illustrates rationalization of a self-benefitting kind.

R2.1. Revision: Distinguishing two functions of rationalization

The pejorative model of rationalization currently dominates. It proposes that while rationalization may have some benefits, improving reasoning is not among them. Not surprisingly, experiments designed to show examples of rationalization in the pejorative sense tend to find that it is associated with bad outcomes. The target article was not intended to deny the existence of self-serving rationalization. It exists. We've documented it. We understand it pretty well. Rather, the contribution of the main article is to point out that the same mechanism – or a very similar one – can also improve reasoning.

Consider an analogy to stereotyping. Prototypically, stereotypes are bad things that make us worse off. This occurs in part because stereotypes can become a vehicle for motivated reasoning, allowing us to distort our factual understanding of the world in order to serve our social interests. Yet, everybody also recognizes that there is also a “rational” side to stereotypes. Given that our minds must make the best use of limited data and limited computation, it makes sense to understand and predict the behavior of tokens (e.g., individual people) by drawing on statistical generalizations at the level of the type (e.g., social group). A common form of computation and representation at the mechanistic level – the stereotype – can serve two quite different functions. One is generally helpful; another can be extraordinarily harmful.

A key outstanding question is whether, and to what extent, self-serving and self-benefitting rationalization share a common mechanism. There is a spectrum of possibilities. At one extreme, it may be that fully distinct mechanisms engage in rationalization of self-serving and self-

benefitting kinds, and that the only connection between these mechanisms is in the abstract form they take. On this view, they are as distinct as the umbrella and the parasol sitting beside each other in a closet – although similarly constructed, they are both physically and functionally distinct. Or, it may be that there is a single, common mechanism that happens to serve two independent purposes. On this view, rationalization is like a roof: A single object that serves the dual purposes of sheltering from rain and from the sun.

R3. Does rationalization really happen?

Tierney & Uhlmann, Stafford, Dahl & Waltzer, and Pärnamets, Johansson, & Hall

(Pärnamets et al.) argue that rationalization is a rare and weak force in our minds. These concerns must be taken seriously because they are supported by some of the very evidence cited in the target article. In a study by Sharot et al. (2010), for instance, the experimentally induced perception of having chosen one vacation destination over another brought about only a 0.06 point shift in subsequent preference, on an 6-point scale.¹ Although there are some published

¹ In the target article I wrote of this finding: “These cases are hard to see as anything but gross errors.” Some of the commentaries interpreted this to mean a *large* error, but it was intended instead in the sense an *undeniable* error. (See, e.g., the first entry of the Merriam Webster Online definition of “gross”: “glaringly noticeable usually because of inexcusable badness or objectionableness // ‘a gross error.’”) I also wrote that, “Sensibly or not, people rationalize all the time.” Some of the commentaries interpreted this rather literally, almost as if I had meant, “At most points in time, people are rationalizing.” Rather, I intended to convey, “People rationalize a lot more than you might have thought, given that it is typically considered ‘not sensible.’” By analogy, “The president plays golf all the time” is meant to convey that she plays a lot more golf than one might have supposed, given the demands of her office.

large effect sizes of rationalization, these commentaries make a persuasive case that the effects of standard laboratory methods (such as the “free choice paradigm”) are usually quite small.

It is notable, however, that standard experimental methods of eliciting rationalization focus specifically on cases where it would be *irrational*. In other words, social psychologists tend not to induce rationalization by prompting people to act in an adaptive manner and then, through a bit of self-reflection, drawing useful lessons from it. Instead, they tend to devise circumstances in which the experimenter has cleverly induced somebody to act in a random or even maladaptive manner, and thus any rationalization would necessarily corrupt reasoning.

This is a useful approach from the standpoint of experimental design. In order to draw strong inferences about psychological mechanisms, it is helpful to put a participant’s behavior under experimental control and then demonstrate that they do something predicted by the mechanism, but otherwise irrational and, therefore, very difficult to explain in any other way (see Saxe 2005).

Unfortunately, the very same features of an experiment that allow us to draw strong inferences may tend to generate weak effects (Mook 1983). If rationalization “approximate(s) a well-understood form of rational inference” (target article, sect. 3.3, para. 3) then, when people are tricked into believing they chose Greece over Thailand, they may rationally infer little from it. Rather, they may conclude that this “choice” was mostly a matter of chance (Gershman 2019; Hawthorne-Madell & Goodman 2019). And, therefore, they may exhibit little change in their beliefs or attitudes.

To see why, it helps to focus on the details of this particular experiment: Although exquisitely designed for strong causal inference, it is wholly unlike ordinary life. Participants are instructed to click on one of two nonsense words that, they are told, contain subliminal information eliciting an unconscious decision between vacation destinations. In fact, the experimenter

randomly assigns the participant one vacation destination or the other. Upon “learning” that their “choice” was Greece (for instance), the participant might begin to introspect a bit. “How do I feel about Greece? Do I really have strong feelings that it is a better vacation destination?” Due to the very logic of this experiment, on average the answer to this question must be “no,” since the participant’s apparent choice is randomly assigned. In light of the peculiar circumstances of the choice, and lacking any actual underlying attitude supporting it, people may reasonably conclude that there are few “true beliefs” or “useful desires” to be gleaned. Evidence for rationalization in this context allows for a strong causal inference, but we ought to expect the effect size to be very small.

Contrast this with a person who is *actually* choosing between vacations in Greece and Thailand – say, in a travel agency, with colorful brochures and detailed information. They feel a strong, intuitive pull toward Greece, but can’t yet put their finger on why. They try to explain their feelings to themselves and to the travel agent. In this circumstance, they might reasonably conclude that there is something true, and something important, to be learned during self-reflection.

Dahl & Waltzer raise a second objection to the thesis that rationalization makes much contribution to human thought and behavior. They write, “If rationalization were widespread, we would have to abandon a premise of much discourse: that our beliefs and actions are generally based on reasons.” Along similar lines, **Tierney & Uhlmann** summarize findings that “Time 1 explicit attitudes predict Time 2 behaviors far better than past behaviors predict future self-reported attitudes, calling into question the prevalence of post hoc rationalizations for past actions... [and] relegating the ‘rationalizations are rational thesis’ to address only a small portion of the attitude-behavior relationship.”

These objections depend on the premise that rationalization is opposed to reason, and so the more influential one is, the less the other must be. Yet, the very premise of the target article is that rationalization can actually improve reasoning. Indeed, the theory depends upon the premise that we are often rational creatures whose Time 1 beliefs and desires appropriately shape our Time 2 behaviors; after all, there is no point exchanging representations in order to improve reasoning if it is impotent or inert! (See sect. R5, where the target article is defended against the critique that we never actually reason).

By analogy, consider the proposition: “Humans frequently sleep in order to benefit their waking lives.” It hardly counts as an objection to say, “But if we sleep so frequently, this leaves no time for us to be awake!” Likewise, “Research suggests we accomplish nearly all tasks while awake, relegating to sleep to an inconsequential role!” Just as a few hours of sleep render our many waking hours more productive, so too might a few moments of rationalization support many moments of reasoning.

Berthelette & Kalbach agree that nearly all adaptive behavior traces back to reasoning, but they arrive there by a different route. They argue that when habits work well – that is, when they are appropriately attuned to circumstances and recommend the correct behavioral policy – they are “best explained as being caused by a standing belief and desire” Thus, there can be no rationalization “because the relevant beliefs and desires were there all along.” The implied psychological model seems to be that habits arise exclusively through the stamping-in of rational actions through repetition (something like practicing a piano piece until it becomes “muscle memory”). This kind of habit formation is well documented (Dezfouli & Balleine 2012; 2013; Dezfouli et al. 2014; Miller et al. 2019). But it is also well-documented that we cache the context-dependent value of actions based on a history of reward (Schultz et al. 1997; Glimcher

2011; Morris & Cushman 2019). In such cases the habitual action does not depend, either presently or historically, on a generative causal model – that is, on beliefs (Dolan & Dayan 2013). In other words, habits are not always just “precompiled rational acts.” Rather, they are sometimes formed by a computationally and representationally distinct process.

In summary, the objection that “rationalization doesn’t happen” begins with the premise that most rationalization is self-serving. Standard social psychology paradigms are well designed to estimate the effect of this “bad” kind of rationalization, and to the extent that it occurs, it stands in opposition to sound reasoning. In contrast, however, the target article proposes a second kind of rationalization. Its influence may be systematically underestimated by standard paradigms, and it is not opposed to reason. Future experiments should explore whether, how often, and how powerfully it shapes our thinking.

R4. The mechanics of rationalization: Prioritized representational exchange

Several commentaries focused on the mechanisms underlying rationalization and representational exchange (**Harmon-Jones & Harmon-Jones, Sievers, Simon & Holyoak**).

The target article focused on a simple case: The post hoc rationalization of action. Here, a person has already performed an action and adjusted their beliefs and desires in a manner that would have rendered the action rational.

The commentaries offer two basic extensions of this model. First, rationalization is often antecedent: “we rationalize our actions not only after we perform them, but also before we perform them and sometimes as a condition of performing them” (**Ellis & Schwitzgebel**). Similarly, we often rationalize not an action, but an attitude or emotion (**Railton**). Second,

rationalization is not merely a process of updating beliefs and desires conditional on action, but a more general process of achieve coherence among a variety of mental state representations, with multidirectional influences.

R4.1. Antecedent rationalization

Sievers and **Simon & Holyoak** assemble impressive evidence for antecedent rationalization, in which a person adjusts their beliefs and desires prior to action (see also **D’Cruz, Ellis & Schwitzgebel**, and **Tierney & Uhlmann**). As Sievers points out, “Nothing entailed by antecedent rationalization precludes the correctness of Cushman’s model of rationalization as a representational exchange mechanism. It does, however, demand an expansion of that model.” This is an important point: Antecedent rationalization happens. It fosters representational exchange when the anticipated action is influenced by a non-rational adaptive control mechanism (e.g., habit, norm, or instinct).

In a paradigmatic case of antecedent rationalization, we bring beliefs and desires into alignment with a non-rational system, such as habit. For instance, suppose that cached value representations (habit) recommend choosing cake instead of tea for dessert, but that a planning system (reason) recommends tea over cake. We would adjust our reasons for tea (“I’ve already eaten too many calories in this meal...”) so that they become aligned with our habitually favored choice (“...but it’s Lisa’s birthday, and she’d be disappointed if I didn’t have some cake”). In this case our cached values remain constant while our reasons change.

Extending this idea, **Railton** proposes that we often rationalize our emotions, and not just our actions. Thus, for instance, suppose that a coworker inspires intense feelings—perhaps admiration, jealousy, anger, love or contempt. Even if these feelings do not move us to any particular action, they still may shape our beliefs by rationalization. In this case, we are learning

what to think or want not by observing how we act, but instead by observing how we feel. In this sense it is akin to antecedent rationalization, which is also prompted by a thought, rather than an act.

But once we acknowledge that certain non-rational forms of thought (e.g., emotion, or habit) can antecedently adjust the representations that contribute to reasoning (beliefs and desires), surely we must also acknowledge that the influence could also run the opposite way, as well. To return to the example above, couldn't the reasons remain the same, while the habit-based values are adjusted so that tea, rather than cake, becomes the habitual response (e.g., Gershman et al. 2014)? These are both instances of representational exchange, but they push information in opposite directions. In one case, information flows from habit to reason; in the other case, from reason to habit.

Presumably, the controller that you “trust” more – the one that is most likely to generate the correct policy – should both determine your action and serve as the source for representational exchange. Deciding which controller to trust is often called the *metacontrol* problem. Classically, metacontrol is construed as a process of deciding which system governs action. (In this manner, it is pivotal to post hoc rationalization, which depends upon the control system underlying action). Importantly, however, it could also be repurposed as a mechanism for determining the direction of representational exchange, for instance, during antecedent rationalization. Metacontrol determines which system should be trusted, and it is an active area of research with several promising theoretical approaches in development (e.g., Daw et al. 2005; Griffiths et al. 2015; Kool et al. 2017; Shenhav et al. 2013;).

In summary, antecedent rationalization surely occurs. But it draws our attention to a crucial problem faced by the theory of representational exchange: How is the direction of exchange

established? The existing literature on metacontrol, which is similarly concerned with the problem of deciding which control systems to trust, offers promising leads. It also immediately suggests a more dynamic model in which rational and non-rational processes continually exert influence upon each other in a manner prioritized by metacontrol.

R4.2. Coherence

Simon & Holyoak and **Harmon-Jones & Harmon-Jones** both contrast the theory of representational exchange with alternative models focused on coherence. To quote Simon & Holyoak, “In models of *coherence-based reasoning*... complex decision situations are represented by networks in which the relevant variables are interconnected via excitatory and inhibitory links. Constraint satisfaction mechanisms settle the network into a stable state of coherence, in which mutually supportive connections (i.e., those that “go together”) activate one another and collectively inhibit their rivals” (emphasis added).

Simon & Holyoak’s commentary has a lot to say about the mechanisms by which coherence is achieved, but little to say about why coherence itself is a useful property of the mind. Several other commentaries furnish possible adaptive explanations. **Harmon-Jones & Harmon-Jones** appeal to the idea that coherence reduces internal conflict: “The organism will have a difficult time enacting a course of behavior when trying to follow the action tendencies evoked by conflicting cognitions.” And we have already summarized several related proposals that coherence has the useful social function of rendering our behavior more predictable to others (**Levy; Saxe & Nettle; Stanford, Thomas, & Sarnecka [Stanford et al.]**). Notably, however, these explanations focus on intrinsic benefits of coherence: that is, the lack of discrepancy or conflict between representations. Rather than positing that we prioritize “better” representations and adjust “worse” ones, the default assumption seems to be that conflicting representations will

exert equal and opposite influences upon each other.

It is possible, of course, that people are designed to achieve coherence among their mental states merely for the sake of coherence itself, without attempting to increase the average quality of the resulting representations. But as long as our minds are achieving coherence, shouldn't they hold relatively fixed those mental state representations that are most likely to be adaptive, while adjusting those that are most likely to be maladaptive? By analogy, consider a marketing company with one headquarters in New York and another in L.A. It would make the company more efficient and nimble to consolidate into a single headquarters. One possibility is to split the difference and relocate to Topeka – at least it would increase coherence! Another possibility, however, would be to carefully consider the performance of the New York and L.A. offices, and then to move the weaker headquarters to the location of the stronger one. This solution achieves coherence with a much more promising outlook for the bottom line. But it also makes greater demands on the decision maker. Like antecedent rationalization, it requires a method of prioritizing some representational exchanges over others.

R4.3. Revision: Coherence with priority

These commentaries offer two important amendments to the original theory of representational exchange: One mechanistic, and one functional. The mechanistic amendment is simply to point out that rationalization is not always ex post, but may often be ex ante, or a dynamic process in which coherence is achieved among many sorts of representations at any particular time. The functional amendment addresses the question at the heart of the target article: Why bother with rationalization? The target article offered two reasons; the commentaries suggest a third.

First, the target article suggested that we rationalize in order to improve reasoning, imparting truer beliefs and more useful (i.e., fit) desires. But this answer must be incomplete, because

natural selection does not directly favor superior reasoning for its own sake; it only favors superior behavior. Why would we expect rationalization to improve the overall behavior of the organism (see **Berthelette & Kalbach**)? To the extent that the organism is already acting well based on some non-rational system, what use is there for improvement to its beliefs and desires?

The target article addressed this question in the broader context of representational exchange, offering a second kind of explanation. Different ways of representing information have different costs and benefits. Among the distinctive benefits of reasoning are its flexibility (i.e., the ability to rapidly adjust in light of new information) and its unique form of generalization (i.e., the ability to generate a policy in novel circumstances according to underlying causal principles). Thus, extracting information from non-rational systems may enhance an organism's capacity for flexibility and generalization, improving its behavior overall. Meanwhile, other directions of representational exchange (e.g., from reasoning to habit, as during habitization) can leverage the distinct advantages of other forms of behavioral control (e.g., the speed and cognitive efficiency of habitual action).

The commentaries, however, draw our attention to another manner in which representational exchange can improve our behavior. If it can successfully prioritize from moment to moment which behavioral systems get to be the source versus the receiver of representational exchanges, this would ensure that representational exchange would lead not only to coherence among systems, but also improvement among them. This stands out as a key area of development for the theory of representational exchange. Existing theories of metacontrol provide a natural starting point.

R5. Social functions of rationalization

The theory of representational exchange proposes an “ultimate” or “computational”-level account of rationalization (Marr 1982; Tinbergen 1963). In other words, it is supposed to answer the question, “Why would natural selection favored rationalization?” Several commentaries also take up this question. Some extend the theory of representational exchange, while others propose alternatives to it. Nearly all of these contributions, however, focus on social functions of rationalization.

R5.1. Cohesion

Van Bavel, Sternisko, Harris, & Robertson (Van Bavel et al.), and Gelpi, Cunningham, & Buchsbaum (Gelpi et al.) propose that rationalization promotes social cohesion. When we rationalize the others’ behavior, or collective cultural practices, we bring our beliefs and desires into closer alignment with that of others. The result is cohesion and a sense of shared identity, which, they propose, has direct personal benefits. As Gelpi et al. write, “Sharing a belief with those in one’s community is therefore beneficial not only when (and because) that belief is true, but also when (and because) it provides an individual with the benefits of a group, such as a sense of belonging and easily accessible shared knowledge.”

Certain versions of this thesis are closely related to the persuasion and blame-avoidance functions of rationalization discussed in the target article (**Altay & Mercier**, see also Haidt 2001; Mercier & Sperber 2011; Tedeschi et al. 1971; Tetlock 2002; Von Hippel & Trivers, 2011). For instance, suppose a political leader is caught performing an apparently corrupt action. In this case she and her supporters might rationalize her behavior by attempting to provide benign explanations for it (Altay & Mercier). But the commentaries also consider extensions of this thesis to many other kinds of cases: We might rationalize collective practices like holidays and sports; shared norms of how we spend time, organize families, and maintain friendships; and,

perhaps above all, moral values (**Gelpi et al.** and **Graham**).

Van Bavel et al. argue that, when it supports social cohesion, “rationalization often facilitates distortion of rather than approximation to truth.” **Altay & Mercier** and **Gelpi et al.** mostly share this view. In other words, theirs is an account of self-serving, pejorative rationalization (albeit one that serves an important social purpose). They do not consider it maladaptive, however, because it pays to get along with one’s peers. (In contrast, **Graham** and **Laurin & Jettinghoff** raise the possibility that rationalizing collective action can sometimes be adaptive due to its self-benefitting nature. This possibility is considered at length in sect. R7.2.)

R5.2. Coordination

Saxe & Nettle, **Stanford et al.**, and **Levy** propose an alternative social function of rationalization: Facilitating coordinated joint action. Humans are nature’s paramount cooperators; we have an unmatched ability to flexibly coordinate our behavior to accomplish joint goals. We succeed in part because we can predict each other’s behavior. We mostly predict each other’s behavior by assuming rational planning and choice (Baker et al. 2009; Dennett 1989). Thus, the commentators propose, it will be easier to explain and predict another person’s action if it exhibits consistent adherence to norms of rationality across time. Rationalization may be useful because it imposes such consistency. As Levy writes, “By making ourselves predictable, we enable more efficient cooperation, which is essential for the flourishing of social animals like us (Tomasello 2014). If each of us can predict how the others will behave, we can more efficiently play our part in joint actions, without interfering with one another or introducing redundancies.”

If anything, this understates the case. Successful joint action often requires not only that we model each other’s intentions, but also that we adopt joint intentions, in which the behavior of

each partner is generated by shared goals and common knowledge (Kleiman-Weiner et al. 2016; Tomasello 2005). Possibly, then, rationalization allows us to translate the behavioral prescriptions of non-rational systems into a format amenable to joint intentional action.

R5.3. Communication

Levy and Pärnamets et al. consider a final social function of rationalization: to foster communication between people. Whether or not reasoning is the language of thought, it is certainly the language of language. We explain ourselves with words, and these words typically express our beliefs, desires, plans, and reasons. Perhaps this is because beliefs, desires, plans, and reasons are especially effective kinds of representations for communication.

It is obvious, for instance, that one person cannot describe their instinct (i.e., innate behavior) to another person and thereby cause the other person to acquire an *instinct*. Rather, some change of representational format is required. In theory, a person could express their instinct (or habit, or norm) simply as a behavioral prescription. But it will often be much more useful to learn the reasons behind a person's policy than to simply learn the policy itself and be forced to copy it. This is because each of us occupies different circumstances, and so the optimal policy for one person may be suboptimal for another. Suppose, for instance, you love peanuts, but I am allergic to them. If you tell me, "Eat these cookies!" I may be substantially worse off than if you tell me, "I eat these cookies because I love the peanuts in them." Rationalizing prior to communication presents information in a format such that a person can flexibly adapt socially learned information to their own unique circumstances.

The target article contrasted two possible pathways of social learning (see Fig. 3 of the target article). In the first, a person acts, an observer engages in theory of mind to extract the beliefs

and desires implicit in the action, and the observer finally adopts those beliefs and desires herself (i.e., by informational conformity). In the second, a person acts, an observer directly copies the behavior (i.e., by normative conformity) and, upon rationalizing her own behavior, the observer finally extracts its implicit beliefs and desires and adopts them. The key insight of **Levy and Pärnamets et al.** is to add a third pathway: A person acts, he rationalizes his own behavior (extracting beliefs and desires from it), and he verbally reports these to an observer who then adopts them. This is a promising addition to the framework.

R5.4. Revision: The social dimension of representational exchange

Each of these reviews posits that rationalization has important and diverse social benefits, fostering cohesion, coordination, and communication. Some of these functions, such as group cohesion, share little overlap with the theory of representational exchange. They more likely explain self-serving rationalization. Others of these, such as coordination and communication, are likely explanations of self-benefitting rationalization. They show that effective representational exchange improves not just our personal thoughts and actions, but also our collective ones.

What is most remarkable about this set of commentaries, however, is their implied commitment to a bold thesis: That social benefits are not just an important explanation of rationalization, but also of reason itself. In fact, there is an even more extreme version of this hypothesis, according to which reason (beliefs, desires, plans, etc.) plays no important role in guiding our behavior, but is instead a socially useful way of talking about our behavior. The next section considers several proposals that engage this possibility.

R6. Do we reason at all?

The target article argued that often the function of rationalization is to improve reasoning by improving its raw materials: The beliefs and desires that it assembles into plans. **Levy, Roskies,** and **Veit, Dewhurst, Dolega, Jones, Stanley, Frankish, & Dennett (Veit et al.)** question whether humans actually reason – in other words, whether we hold beliefs and desires and use them to make plans. None of these authors dispute that we talk in terms of reasons, beliefs, desires, and plans. And, they agree that such talk has meaning and utility. But, they suggest, maybe it's just talk. (Some endorse this possibility more strongly than others). The target article described rationalization as a “useful fiction,” but is the real fiction reason itself?

According to these critiques, people exhibit “flexible response[s] to environmental and internal information,” but they do so in ways that do not depend on “explicit representations” of belief and desire (**Levy**). Rather, what exist are “some sort of proto-mental states” (**Veit et al.**), the “great majority of these states are dispositional and may never be explicitly represented” (**Levy**). **Roskies** offers a clear introduction to these ideas, employing the helpful example of a thermostat: “To the degree that we can predict and explain a system’s behavior by imputing folk psychological states [e.g., beliefs, desires, plans, and reasons] to that system, so we are warranted in that imputation. For Dennett [1987], even simple systems such as thermostats are legitimate targets of the intentional stance, even though no standard realist would support the view that thermostats have mental states.” The proposed function of the intentional stance, not dissimilar to Bem’s (1972), is not to improve one’s own reasoning, but rather to allow us to describe, predict, and explain our own and others’ actions.

There is much to like in these commentaries, but also something amiss. Surely, it is often useful to impute reasons to psychological systems that do not reason. This is the essence of the

target article's model of rationalization as representational exchange. And, as discussed in section R4, rationalization has benefits not only for the individual who rationalizes, but also for those who depend upon her for collaboration and communication. But here is the key question: Do these facts explain away reasoning itself? If we accept that our experience of having reasoned in sometimes a fiction, do we have any reason to accept that it is ever a reality?

We do. As defined in the target article, to act by reasoning is to estimate the expected value of candidate actions by consulting a generative model of their consequences, and to then choose actions proportional to their expected value (following, e.g., Daw & Dayan 2014). There is ample evidence that people learn and represent generative causal models (Buckner et al. 2008; Gershman 2018; Gläscher et al. 2010) and can use these to assess the likely consequences of candidate actions (Deserno et al. 2015; Doll et al. 2015; Simon & Daw 2011; Vikbladh et al. 2019). By considering these consequences in light of our goals our preferences, we can estimate the instrumental value of candidate actions (Bakkour et al. 2019; Daw et al. 2011; Jones et al. 2012). These value estimates reliably influence choice (Cushman & Morris 2015; Daw et al. 2011; Deserno et al. 2015; Doll et al. 2015; Simon & Daw 2011). These are precisely the kinds of representations and computations that can be improved by representational exchange.

In summary, these commentaries emphasize the striking and illuminating fact that we can profitably understand a thermostat by treating it as rational, even if it is not (i.e., by adopting the intentional stance). This immediately suggests a subversive and alluring possibility: Maybe, despite feeling and talking like rational creatures, we are more like very complicated thermostats. Yet, as an empirical matter, we are not. When a thermostat adjusts the temperature, it is not by encoding a generative model of the world, computing the likely consequences of various actions, and then choosing actions proportional to the resulting instrumental value estimates. But, when a

human adjusts the thermostat, it often is.

R7. Making meaning

We cannot always know the causes of our behavior, but rather must infer them (Bem 1972; Nisbett & Wilson 1977). We do this in part because it brings order and meaning in our lives. Many theories posit that we search for order and meaning because it makes us feel better, which is a proximate explanation (Baumeister 1991; Hasselkus 2011; Reker et al. 1987; Webster & Kruglanski 1994). The theory of representational exchange provides a complementary ultimate explanation: The meaning we make out of our actions, attitudes and emotions (**Railton**) renders useful information available for reasoning. A final set of commentaries takes this idea of meaning-making as its point of departure. The authors observe how phenomena as diverse as daydreaming (“everyday thought experiments”: **Ullman**), system justification (**Laurin & Jettinghoff**), and ideological commitment (**Graham**) can be understood as attempts to find meaning, breaking new ground for the theory of representational exchange. They make a convincing case that these are powerful and constant forces in our thinking. They also offer a cautionary note, questioning what happens when even self-benefitting rationalization is misapplied (**Weinberg & Weinberg, Pärnamets et al., Graham**).

R7.1. Rationalization and personal meaning

People daydream about all kinds of fantastical things. They love it. They’ll do it on their own, and as **Ullman** shows, they’ll happy to do it if you ask. This is remarkable because many of the questions posed by Ullman et al. (2016) are patently outlandish:

Imagine that aliens come down to Earth, and give you the option to go with them

on their travels throughout the universe. The aliens are friendly and honest, and tell you that you would see amazing things on your travels with them if you decide to go with them.

If you decide to go, you will have a week to say goodbye to your family and friends. Once you leave, you will never again return to Earth, nor be able to communicate with people on Earth.

Do you go?

Different people give different answers, and people are bad at guessing what others will say, yet most people report being highly confident in their own personal answers. Not only do people say they enjoy thinking about things like this, they say that they *learned* something from it – mostly things about themselves:

That I am more risk averse when faced with transformative decisions.

That I am attached to my actual, earthly existence more than I might have thought otherwise.

That I value personal relationships more than I thought.

These findings comport with a large literature showing that when people have downtime, their minds turn toward thoughts about imagined possibilities for themselves, in the past and in the future (Buckner et al. 2008).

As **Ullman** notes, the theory of representational exchange “can help explain why such everyday thought experiments are informative, and also why people like to engage in them.”

When people surprise themselves with their own confident answers to hypothetical questions, it means that they are discovering implicit information in their automatic reactions, including their emotions (**Railton**) – information that often contradicts their explicitly held beliefs and values. They consider this information valuable and *true* – they take it to reflect important insights about who they are and how they ought to feel. By organizing our impulses around meaning, purpose, and narrative, we unpack implicit information from non-rational systems and make it available to improve subsequent reasoning.

R7.2. Rationalization and social meaning

Although the target article focused mostly on the rationalization of one's own actions, **Graham** and **Laurin & Jettinghoff** propose a key extension, and Graham writes: “Processes like rationalization often occur at the collective levels of groups, societies, and cultures. Collective rationalization – whereby a group's collective action leads them to update or create new shared beliefs and desires – may be rational as well. Just as individual rationalization can extract useful information from non-rational sources like instincts and habits, communal processes can cohere moral intuitions and norms into the shared ‘useful fictions’ of shared moral narratives.”

Graham's commentary focuses on ideology: A shared sense of meaning – an interconnected set of beliefs and values that give a sense of coherence, explanation and purpose to groups (Jost 2006). Laurin & Jettinghoff's commentary focuses on system justification: Our tendency to assume that prevailing social structures are good. They agree that we seek not only to ascribe meaning to our individual actions, but also to collective cultural practices.

Laurin & Jettinghoff wonder, however, why we would be justified in not just inferring but actually adopting the beliefs and desires that we attribute to others. In other words, why drink the collective Kool-Aid? “If people were merely trying to identify beliefs and desires” that make

sense of other's actions, "there would be no reason to expect system justifying rationales to predominate" over system-undermining ones. Why assume the system is good?

Different commentaries suggest different potential answers to this problem. **Gelpi et al.** appeal to cognitive efficiency: "Trusting the beliefs shared by one's social group – and 'outsourcing' one's own cognition to depend on knowledge held by others in their community – can also reduce the need to engage in cognitively effortful reasoning ... (Sloman & Rabb 2016)." In contrast, **Laurin & Jettinghoff** consider the possibility that we do not rationalize the system itself, but rather our own apparent acceptance of it. Perhaps, then, system justification is really justification of one's own complacency. Both of these explanations probably capture a good part of the truth, along with the possibility that the rationalization of collective behavior fosters social cohesion (**Van Bavel et al.** and Gelpi et al.).

But there is also a far simpler answer: We rationalize collective cultural practices, adopting their implied beliefs and values, because collective cultural practices are usually pretty good. Just as genes are subject to genetic evolution, collective practices are subject to cultural evolution (Boyd et al. 2011). Thus, if you observe that most people are doing something (or that the cultural system is organized in some way), it is more likely to be adaptive than maladaptive. This does not imply that the current norm is the best that could possibly exist, any more than we should suppose that a sparrow's wings are the best wings that a sparrow could possibly have. Rather, it implies that adhering to the norm is probably better than ignoring it, just as the sparrow is better off with her wings than without them.

Confronted with a claim like this – *if most people do it, it's probably adaptive* – our minds naturally leap to counterexamples. Wasn't it for good reason that our mothers scolded, "If everybody jumped off a bridge, would you?" After all, people do all sorts of terrible things. Most

people eat too much, don't get enough exercise, are cognitively and socially biased, and so forth. There was a time and a place when most people thought the Earth was flat; when most people smoked cigarettes; and when most people thought it was acceptable to whip your child, your wife, and your slave. Human progress has always relied on those who say, "I can do better than what most people are doing," and it always will.

These counterexamples are important and inescapable. But now imagine what it would mean to abandon our cultural inheritance entirely. Imagine a child who grows up without any inclination to think or act like others. Given soup, she would have no inclination to use a spoon just because her big sister does; rather, she would *reason* about whether a spoon works for her. The same goes for forks, manners, values, laws, school, sports, holidays, fashions, jobs, and so on. Shown a recipe for bread, she would not assume there was anything useful about the yeast, the kneading, the rise, or even the oven. She would just rely on her own knowledge and sound reasoning to deduce the proper manner of baking. Now, whose loaf would you rather eat: Hers or an obedient apprentice's?

Our collective cultural practices are a recipe for life. It is not a perfect recipe – in fact, you can almost certainly improve it with some tweaks. But it is unlikely that you will do better by ignoring the recipe altogether. More likely, any improvements you make will necessarily build upon the cumulative wisdom of generations of cooks that have come before you. (In Newton's terms, seeing further than giants requires not just extending your legs, but also climbing on their shoulders). Sometimes most people do the wrong thing, but these counterexamples miss the forest for the trees – nearly everything we do right was learned, in part, from those around us (Boyd et al. 2011). This accords with a basic impulse of conservatism in political theory and philosophy (e.g., Hayek 1973): Tradition may require revision or rejection, but it at least

deserves respect.

Or, at least, it does if you are interested in fitness within the framework of biological and cultural evolution. As scientists attempting to understand the structure of the human mind – in particular, our propensity to rationalize both personal and collective actions – this is the kind of fitness we are interested in as a descriptive theory. But what was fit for the past may be unfit for the present, and what is fit for nature may be unfit for ethics.

R7.3. Spoiled Kool-Aid: Is rationalization outdated?

While accepting that rationalization may have an adaptive rationale, **Weinberg & Weinberg**, **Pärnamets et al.**, and **Graham** ask: Is it the right thing to do? They offer three reasons to think it is not.

The first is familiar: At least some rationalization is self-serving, and it will generally lead us away from the truth. While it may have some adaptive benefits when it works properly (such as cohesion, as noted by **Van Bavel et al.**), **Graham** documents the pernicious effects when it goes awry, such as the justification of prejudice or even violence toward outgroups.

Graham goes further, however, by questioning even self-benefitting rationalization. He accepts that by adopting the values implied by instincts and norms, we are adopting values shaped by adaptive processes – biological and cultural evolution, respectively. These may be fit, he notes, but are they morally right? This is an important question; indeed, it is among the most important questions we face. If rationalization captures the conservative impulse to respect traditions, we must confront the fact that many human traditions are morally abhorrent. So, too, is blind respect for them.

Weinberg & Weinberg and **Pärnamets et al.** offer a final reason to distrust self-benefitting rationalization: While it may have been adaptive when it evolved, perhaps it is poorly suited to

present circumstances. We evolved at a time when the world was relatively more static. Thus, norms, instincts, and habits shaped by the past were usually adaptive in the present. And, in turn, implicit beliefs and desires extracted from these sources were generally useful. Yet, today the world changes at a faster pace environmentally, technologically, and socially. This drives our instincts, norms, and habits out of date more quickly than ever before. (And while the commentaries do not mention it, it is also likely that our beliefs are truer, and capacity for reason keener, than ever before.) If natural selection struck a balance between non-rational systems and reasoning when rationalization evolved, perhaps it is the wrong balancing point today. This point is worth careful consideration. If indeed the balancing point of representational exchange is innate and shaped by biological natural selection, then it is hard to imagine that it is optimal today. Yet, as discussed in section R4, the balancing act between reason and intuition may be carried out online by metacontrol processes that are relatively more sensitive to present circumstances.

R7.4. Revision: The scope and challenge of representational exchange

This final group of commentaries offers two important revisions to the target article. First, they broaden the reach of the representational exchange to personal and collective self-understanding: The attempt to find meaning in our lives. The target article talked a lot about extracting information, but not at all about making meaning. Yet, the search for narrative, structure, and meaning in our lives is a perfect example of representational exchange. It may be a fiction to suppose that our individual or collective behavior already *has* meaning, as if organized by an intelligent agent with a specific plan in mind. But if it is a fiction, it is surely a useful one. We stand to learn a tremendous amount by assuming that the way each and all of us do things is, while not perfect, probably very good. Second, however, they sound a cautionary note even for

self-benefitting rationalization. It may be adaptive, but is it good? They provide ample reasons for skepticism.

On their face, these commentaries seem to strike two discordant notes. On the one hand, representational exchange is at the heart of our search for meaning. On the other, it may often be a misguided search. Yet, in fact, these notes combine to play a familiar theme: the familiar struggle to balance reasoning against intuition – head against heart.

R8. The human problem

The Enlightenment bestowed a durable understanding of the human problem: To act rationally and, therefore, by reasoning. Summarizing this neoclassical spirit, Pinker (2018) writes: “If there’s anything the Enlightenment thinkers had in common, it was an insistence that we energetically apply the standard of reason to understanding our world, and not fall back on generators of delusion like faith, dogma, revelation, authority, charisma, mysticism, divination, visions, gut feelings or the hermeneutic parsing of sacred texts.”

Although not our only understanding of the human problem, this is an influential one. It aligns with a strong tradition of psychological research that reifies reasoning, organizing it around two normative standards. First, one ought to establish beliefs by the rational application of Bayesian principles. Second, one ought to choose actions by expected value maximization. Against the background of these standards, the behavioral prescriptions of norms, instincts, and habits are defined as biases, mostly to be understood in light of their systematic deviations from reasoning. If they can be justified at all, it is either on the grounds that they are more computationally efficient (Griffiths et al. 2015) or have indirect social benefits (Mercier & Sperber 2011; Tedeschi et al. 1971; Von Hippel & Trivers 2011). These ancillary benefits may compensate for what are otherwise strictly worse outcomes: mistaken beliefs, misguided desires,

and maladaptive choices.

Thus, according to the strictest Enlightenment ideals, we must purge thought and action of the influence of non-rational systems. Pinker (2018), again: “Thinkers such as Kant, Spinoza, Hobbs, Hume and Smith were inquisitive psychologists and all too aware of our irrational passions and foibles. They insisted that it was only by calling out the common sources of folly that we could hope to overcome them. The deliberate application of reason was necessary precisely because our common habits of thought are not particularly reasonable.” This is why we might characterize reasoning not as the human condition, but rather the human problem. It is a problem because, while reasoning is presumed to be superior, it also hard to do. If our highest human calling is self-improvement, then we must apply the force of will to achieve the purity of reasoned thinking.

In its purest form, this Enlightenment vision possesses the allure of simplicity, the venire of logic, and the tragic myopia of self-certainty. Human reasoning is, after all, exquisitely fallible. Our beliefs are often incorrect and always incomplete. Our desires may internally inconsistent, interpersonally irreconcilable, and occasionally detrimental to our well-being. And our capacity for rational planning is severely constrained by time and cognitive bandwidth. None of this implies that reason is bad; to the contrary, it is indispensable. But it is fallible, and sometimes it can be improved by attending to non-rational adaptive mechanisms. Whether through the genetic inheritance of our instincts, the cultural inheritance of our norms, or the experiential inheritance of our habits, we are endowed with guides to action shaped by powerful forces of trial-and-error adaptation.

Thus, to act truly rationally cannot be to act only by reasoning. Rather, even if we accept the Enlightenment commitment to rational action – *especially* if we accept it – then the basic human

problem is not to isolate reason from non-rational processes, but to improve it by considering them. This, of course, is Jason Bourne's problem: To glean from his programmed abilities whatever information serves his present interests. In this effort, rationalization is an essential tool. It is not a matter base self-enhancement or even one of elevated self-understanding. It is, at heart, a project of self-improvement – surely the highest human calling of all.

References

- Baker, C. L., Saxe, R. & Tenenbaum, J. B. (2009) Action understanding as inverse planning. *Cognition* 113(3):329–49.
- Bakkour, A., Palombo, D. J., Zylberberg, A., Kang, Y. H., Reid, A., Verfaellie, M., ... & Shohamy, D. (2019) The hippocampus supports deliberation during value-based decisions. *eLife* 8.
- Baumeister, R. F. (1991) *Meanings of life*. Guilford Press.
- Bem, D. J. (1972) Self-perception theory. In: *Advances in experimental social psychology*, vol. 6, pp. 1–62. Elsevier.
- Boyd, R., Richerson, P. J. & Henrich, J. (2011) The cultural niche: Why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences USA* 108:10918–25.
- Buckner, R. L., Andrews-Hanna, J. R. & Schacter, D. L. (2008) The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences*. 1124:1–38.
- Cushman, F. A. & Morris, A. (2015) Habitual control of goal selection in humans. *Proceedings of the National Academy of Science USA* 112(45):13817–22.

- Daw, N. D. & Dayan, P. (2014) The algorithmic anatomy of model-based evaluation. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369(1655):20130478.
- Daw, N. D., Gershman, S. J., Seymour, B., Dayan, P. & Dolan, R. J. (2011) Model-based influences on humans' choices and striatal prediction errors. *Neuron* 69(6):1204–15.
- Daw, N. D., Niv, Y. & Dayan, P. (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience* 8(12):1704.
- Dennett, D. C. (1987) *The intentional stance*. MIT Press.
- Dennett, D. C. (1989) *The intentional stance* (Reprint edition). MIT press.
- Deserno, L., et al. (2015) Ventral striatal dopamine reflects behavioral and neural signatures of model-based control during sequential decision making. *Proceedings of the National Academy of Sciences USA* 112(5):1595–1600.
- Dezfouli, A. & Balleine, B. W. (2012) Habits, action sequences and reinforcement learning. *European Journal of Neuroscience* 35(7):1036–51.
- Dezfouli, A. & Balleine, B. W. (2013) Actions, action sequences and habits: evidence that goal-directed and habitual action control are hierarchically organized. *PLOS Computational Biology* 9(12):e1003364.
- Dezfouli, A., Lingawi, N. W. & Balleine, B. W. (2014) Habits as action sequences: hierarchical action control and changes in outcome value. *Philosophical Transactions of the Royal Society B: Biological Sciences* 369(1655):20130482.
- Dolan, R. J. & Dayan, P. (2013) Goals and habits in the brain. *Neuron* 80(2):312–25.
- Frederick, S. (2005) Cognitive reflection and decision making. *Journal of Economic Perspectives* 19:24–42.

- Gershman, S. J. (2018) The successor representation: Its computational logic and neural substrates. *Journal of Neuroscience* 38(33):7193–7200.
- Gershman, S. J. (2019) How to never be wrong. *Psychonomic Bulletin and Review* 26(1):13–28.
doi: <https://doi.org/10.3758/s13423-018-1488-8>.
- Gershman, S. J., Markman, A. B. & Otto, A. R. (2014) Retrospective revaluation in sequential decision making: A tale of two systems. *Journal of Experimental Psychology: General* 143(1):182.
- Gläscher, J., Daw, N., Dayan, P. & O’Doherty, J. P. (2010) States versus rewards: Dissociable neural prediction error signals underlying model-based and model-free reinforcement learning. *Neuron* 66:585–95.
- Glimcher, P. W. (2011) Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proceedings of the National Academy of Sciences USA* 108(Suppl. 3):15647–54.
- Griffiths, T. L., Lieder, F. & Goodman, N. D. (2015) Rational use of cognitive resources: Levels of analysis between the computational and the algorithmic. *Topics in Cognitive Science* 7(2):217–29.
- Haidt, J. (2001) The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review* 108(4):814.
- Hasselkus, B. R. (2011) *The meaning of everyday occupation*. Slack.
- Hawthorne-Madell, D. & Goodman, N. D. (2019) Reasoning about social sources to learn from actions and outcomes. *Decision* 6(1):17–60.
- Hayek, F. A. (1973) *Law, legislation and liberty, vol. 1: Rules and order*. University of Chicago Press.

- Heider, F. (2013) *The psychology of interpersonal relations*. Psychology Press.
- Jones, J. L., Esber, G. R., McDannald, M. A., Gruber, A. J., Hernandez, A., Mirenzi, A. & Schoenbaum, G. (2012) Orbitofrontal cortex supports behavior and learning using inferred but not cached values. *Science* 338(6109):953–56.
- Jost, J. T. (2006) The end of the end of ideology. *American Psychologist* 61:651–70.
- Kleiman-Weiner, M., Ho, M. K., Austerweil, J. L., Littman, M. L. & Tenenbaum, J. B. (2016) Coordinate to cooperate or compete: Abstract goals and joint intentions in social interaction. In: *Proceedings of the Cognitive Science Society*.
- Kool, W., Gershman, S. J. & Cushman, F. A. (2017) Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychological Science* 28(9):1321–33.
- Marr, D. (1982) *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman.
- Mercier, H. & Sperber, D. (2011) Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences* 34(2):57–74.
- Miller, K. J., Shenhav, A. & Ludvig, E. A. (2019) Habits without values. *Psychological Review*.
- Mook, D. G. (1983) In defense of external invalidity. *American Psychologist* 38(4):379–87.
- Morris, A. & Cushman, F. (2019) Model-free RL or action sequences? *PsyArxiv* DOI: 10.31234/osf.io/k67tm
- Nisbett, R. E. & Wilson, T. D. (1977) Telling more than we can know: Verbal reports on mental processes. *Psychological Review* 84(3):231.
- Pinker, S. (2018) *Enlightenment now: The case for reason, science, humanism, and progress*. Penguin.

- Reker, G. T., Peacock, E. J. & Wong, P. T. (1987) Meaning and purpose in life and well-being: A life-span perspective. *Journal of Gerontology* 42(1):44–49.
- Saxe, R. (2005) Against simulation: The argument from error. *Trends in Cognitive Sciences* 9:174–79.
- Sharot, T., Velasquez, C. M. & Dolan, R. J. (2010) Do decisions shape preference? Evidence from blind choice. *Psychological Science* 21(9):1231–35.
- Shenhav, A., Botvinick, M. M. & Cohen, J. D. (2013) The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron* 79(2):217–40.
- Schultz, W., Dayan, P. & Montague, P. R. (1997) A neural substrate of prediction and reward. *Science* 275(5306):1593–99.
- Simon, D. A. & Daw, N. D. (2011) Neural correlates of forward planning in a spatial decision task in humans. *Journal of Neuroscience* 31(14):5526–39.
- Sloman, S. A. & Rabb, N. (2016) Your understanding is my understanding: Evidence for a community of knowledge. *Psychological Science* 27(11):1451–60.
- Tedeschi, J. T., Schlenker, B. R. & Bonoma, T. V. (1971) Cognitive dissonance: Private ratiocination or public spectacle? *American Psychologist* 26(8):685.
- Tetlock, P. E. (2002) Social functionalist frameworks for judgment and choice: Intuitive politicians, theologians, and prosecutors. *Psychological Review* 109(3):451.
- Tinbergen, N. (1963) On aims and methods of ethology. *Zeitschrift für Tierpsychologie* 20(4):410–33.
- Tomasello, M. (2005) Understanding and sharing intentions: The origins of cultural cognition. *Behavioral and Brain Sciences* 28:675–35.
- Tomasello, M. (2014) *A natural history of human thinking*. Harvard University Press.

- Ullman, T. D., Siegel, M., Tenenbaum, J. B. & Gershman, S. J. (2016) Coalescing the vapors of human experience into a viable and meaningful comprehension. *Proceedings of the 38th Annual Conference of the Cognitive Science Society*.
- Vikbladh, O. M., Meager, M. R., King, J., Blackmon, K., Devinsky, O., Shohamy, D., ... & Daw, N. D. (2019) Hippocampal contributions to model-based planning and spatial memory. *Neuron* 102(3):683–93.
- Von Hippel, W. & Trivers, R. (2011) Reflections on self-deception. *Behavioral and Brain Sciences* 34(1):41–56.
- Webster, D. M. & Kruglanski, A. W. (1994) Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology* 67(6):1049.