



A Common Framework for Theories of Norm Compliance

Citation

Morris, Adam, and Fiery Cushman. 2018. A Common Framework for Theories of Norm Compliance. *Social Philosophy and Policy* 35, no. 1: 101-127.

Published version

<https://doi.org/10.1017/s0265052518000134>

Link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37373462>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available. Please share how this access benefits you. [Submit a story](#)

A common framework for theories of norm compliance

Adam Morris & Fiery Cushman

*Department of Psychology, Harvard University
Cambridge, MA 02138*

Humans often comply with social norms, but the reasons why are disputed. Here, we unify a variety of influential explanations in a common decision framework, and identify the precise cognitive variables that norms might alter to induce compliance. Specifically, we situate current theories of norm compliance within the reinforcement learning framework, which is widely used to study value-guided learning and decision-making. This framework offers an appealingly precise language to distinguish between theories, highlights the various points of convergence and divergence, and suggests novel ways in which norms might penetrate our psychology.

1. Introduction

Social norms – behavioral prescriptions or proscriptions known by members of a group and informally socially enforced – are central to human societies (Bicchieri, 2006; Cooter, 2000; Fehr & Fischbacher, 2004). They range from fashion trends to religious customs, from norms of honor in the American south (Nisbett & Cohen, 1996) to norms of reciprocity in northeast Uganda (Turnbull, 1987), from honest tax reporting in industrialized countries (Elster, 1989) to conflict resolution among rural farmers (Ellickson, 1994). Thus, it is no surprise that virtually every corner of social and cognitive science has attempted to understand what norms are, when and why they emerge, and how they influence human behavior.¹

It is equally unsurprising that not everybody agrees on the answers. Consider, for instance, the question of how norms influence behavior. People tend to comply with the norms they are exposed to. In many contexts, they choose to adopt the behaviors that the norms prescribe

and avoid behaviors that the norms proscribe. Why? How does exposure to norms influence subsequent decision making? Some prominent theories include: (1) People adopt the statistically dominant patterns of behavior they observe around them (Cialdini & Trost, 1998), (2) people develop heuristic responses to social dilemmas based on their prior history of reward and punishment (Rand et al., 2014), and (3) people learn to derive subjective utility from the absolute and relative welfare of others (Fehr & Schmidt, 2006). As these examples make apparent, contemporary theories of norm compliance are not just saying different things; they're speaking different languages.

Our goal is to propose a *lingua franca* for theories of norm compliance. Our point of entry is a computational framework for learning and decision making that unites key ideas from machine learning, neuroscience, psychology and economics. The framework, known as reinforcement learning (RL), identifies a set of cognitive variables involved in the decision process, any of which norms could intervene on. It offers a precise language in which to enumerate the ways that norms could influence decisions, and specifies exactly how those

¹ For much more nuanced definitions of norms and their kinds, see Bicchieri (2006).

interventions would affect decision makers. In essence, RL is a broad framework that asks how people learn value and then make decisions based on those values, and offers a formal setting in which competing answers can be precisely articulated. It is an ideal framework for theories of norm compliance, which also involves learning and deciding based on representations of value.

We map various influential theories of norm compliance onto this framework, compare their commitments and predictions, and propose novel potential mechanisms of norm compliance. Our analysis suggests that norms influence many parts of the decision-making process, including ones not yet considered, and that norm compliance is grounded in similar computational principles as other forms of human decision making.

In this endeavor, we focus on norms which are intuitively "moral": Norms for fairness, cooperation, reciprocity, etc., which instruct people forgo personal gain to promote the well-being of others. (Think of norms for contributing to the public offering at church, or for recycling, or for offering fair monetary splits in an Ultimatum Game (Thaler, 1988).) Moral norms offer an ideal testing ground for theories of norm compliance, because their influence is strong enough to overpower other competing interests. Which parts of the decision process are altered by exposure to moral norms, and how do these alterations make people more likely to obey the norm?

2. A framework for modeling value-guided decision-making

To answer this, we need a cognitive framework of learning and decision making. We take as our starting point the idea that humans often make decisions by putting their options on a common scale of subjective "goodness", and biasing choice towards options with more of it. In other words, they choose according to each option's value. The idea of value-based choice has a long history in both psychology (see Dolan & Dayan, 2013) and economics (Samuelson, 1948; von Neumann & Morgenstern, 1944), and has

recently assumed an equally prominent role in neuroscience (Glimcher & Fehr, 2013).

Value can be interpreted in at least two ways. One interpretation, common in economics, is that value is simply an abstract description of consistent choice, and not necessarily an actual variable computed by humans during the decision process. However, advances in cognitive and neural science in the last twenty years indicate that humans indeed compute a representation of the overall goodness of an option, and use it to guide their choices (Glimcher & Fehr, 2013). We therefore interpret value as a real cognitive variable computed during the decision process (Becker, 1996; Sen, 1997; Sunstein, 1996).

Our aim is to ground current theories of norm compliance in the RL framework for value-guided decision-making in humans. Of course, this presumes that decisions to comply with norms involve computing the value of options. But this is not a necessary truth; not all behavior is value-guided. One major counterexample is reflexive behavior, which is triggered by perceptual and cognitive processes that do not explicitly represent value in any meaningful sense (Pavlov & Anrep, 1927). Moreover, many non-reflexive choices appear to be driven by processes that do not place options on a common scale, and therefore are not value-guided (Lichtenstein & Slovic, 2006; Vlaev, Chater, Stewart, & Brown, 2011).

So there are viable alternatives to a value-guided theory of norm compliance; indeed, we will return to consider some reasons to take alternative approaches very seriously. One key aim of the computational framework we present below is to make precise what it actually means to say that decision-making is guided by value, such that the claim that norms influence behavior via such representations is both substantive and falsifiable.

With these caveats in mind, however, there are at least two major reasons to begin with the assumption that norm compliance operates in humans principally via processes of value-guided decision-making. First, people's choices to comply with norms show signatures of rational planning over a common currency objective. Their choices in economic games conform to the maxims of consistency (Andreoni, Castillo, & Petrie, 2003; Andreoni & Miller, 2002), and obey the law of demand: As the cost of being fair, reciprocal, etc., increases,

³ We use the terms "obey", "follow", "comply with", etc., interchangeably to indicate situations where people do what a norm prescribes (or avoid doing what a norm forbids). Unlike some authors, we do not use the terms to differentiate hypotheses about the psychology underlying that behavior (Kelman, 1958; Koh, 1997).

people do it less (C. M. Anderson & Putterman, 2006; Andreoni & Vesterlund, 2001; Capraro, Jordan, & Rand, 2014; Carpenter, 2007). This is intuitively true for real-world examples of norm compliance also: People are more likely to cheat on their taxes when it will save them more money, or run a red light when they are in a rush. The ability to flexibly trade off compliance with other goods is a hallmark of value-guided decision making.

Second, prototypically norm-compliant choices are accompanied by increased activity in brain regions involved in domain-general value computation (Ruff & Fehr, 2014). When people choose to allocate money fairly, or punish someone who violated trust, or mutually cooperate in a Prisoner's Dilemma, they show patterns of activation in the ventromedial frontal and orbitofrontal cortex consistent with value-guided choice—even when doing so causes them personal financial loss (Quervain et al., 2004; Rilling et al., 2002; Tabibnia & Lieberman, 2007; Zaki & Mitchell, 2011). These same regions appear to compute and store the value of options across a wide variety of other domains (Glimcher & Fehr, 2013), suggesting that people's decisions to comply with norms are guided by similar processes.

In summary, there is good reason to believe that norm compliance occurs at least in part by processes of value-guided decision-making. Next, we present a powerful formal computational framework for understanding the mechanics of such processes. It comprises two parts: A statement of the decision-making *problem*, which is formalized as a Markov Decision Process (MDP), and a description of the kinds of *algorithmic solutions* that have been pursued to solve that problem, which are called reinforcement learning (RL) approaches.

This framework will suggest five cognitive decision variables that could be influenced by norms. We will map four of these onto existing theories of norm compliance, and offer the other as a direction for future research.

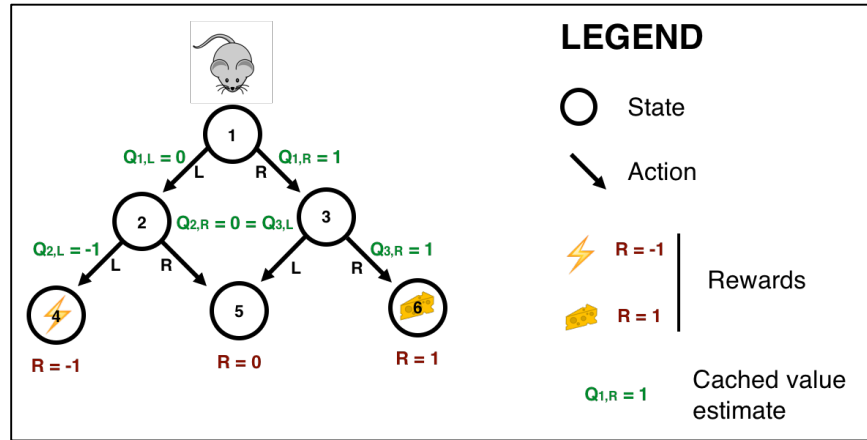


Figure 1: MDP representation of a simple decision problem. A rat must decide how to traverse a maze (represented as a set of 6 states, each with a set of potential actions) to maximize long-term accumulation of reward. R_s represent the intrinsic rewards associated with each terminal state, and Q_s represent the average value (i.e. long-term expected future reward) that the rat would learn to associate with each prior action. For instance, turning right at State 1 eventually leads to the cheese in State 6, and the rat would therefore learn that $Q(\text{state } 1, \text{action } R) = +1$.

2.1 Defining the problem: Markov decision processes

Before defining MDPs in abstract terms, it is helpful to orient ourselves around the concrete details of an idealized example. Imagine a rat, Lenny, in a maze with cheese and electric shocks (Fig. 1). Lenny must decide how to traverse the maze in a way that achieves his objectives. As described above, we assume that the subjective goodness or badness of the cheese and electric shocks can be put on a common numeric scale, which in the MDP framework is called reward. For example, Lenny's rewards might be: $CHEESE = +1$, $SHOCK = -1$. Lenny, then, has a single objective: maximize the long-term accumulation of reward. Lenny could then compute the amount of long-term reward gain expected from turning 'left' or 'right' at each fork in the maze, and simply choose the option with the highest amount.

This situation—an agent has a set of well-defined choices with well-defined consequences, and her goal is to maximize a single numerical objective—is what MDPs formally describe. MDPs carve the world into a set of discrete states S , which can be either decision points or terminal nodes. The decision points each have a set of actions A . (At terminal nodes, there are no more actions to be taken [e.g. the end of the maze], so $A = \emptyset$.) The decision points capture all the moments where the person could do meaningfully different things. For instance, Lenny could carve the maze up into the states in Fig. 1, where each fork is a decision point with the actions "left" and "right".

But note that this carving is subjective: S and A are not objective features of the external environment, but variables in the person's head. Lenny could have incorporated the temperature of the room and made separate decision points for "initial fork in hot room" and "initial fork in cold room". Or he could have codified each physical step as a decision point, with the actions "move left foot forward", "move right foot forward", "move left foot backward", etc. States and actions can also be specified at various levels of abstraction (Cushman & Morris, 2015; Gershman & Niv, 2010). Lenny could conceptualize his actions as "turn left" or "turn right", but he could also conceptualize them as "turn body 27 degrees counterclockwise", "move left foot 1.36 centimeters forward", etc. These variables are subjective and, therefore, could be influenced by norms.

MDPs also formalize the consequences of choice. There are two questions to answer after making a choice: What happened, and how good or bad was it? In MDPs, after selecting an action at a decision point, an agent transitions to a new state (either another decision point, or a terminal node). This process is captured by a noisy transition function $T(s, a, s')$, which gives the probability of transitioning to any new state s' after choosing an action a at a decision point s . This is the answer to the question, "What happened?"

Then, the agent receives a reward. In principle, how much the agent likes or dislikes the consequences of a choice could be based on all three variables: where they were (s), what they chose (a), and where they ended up (s'). Think of a daredevil jumping a canyon on a motorcycle. Her reward from this experience will depend on where she was (at the edge of a canyon), what she did (ride a motorcycle forward), and where she happened to end up (on the other side, or at the bottom of the canyon). Thus, to relate the person's objective to the consequences of choice, we introduce a reward function $R(s, a, s')$, which gives the amount of reward a person receives for taking action a at decision point s and ending up at state s' . Note that, just like the previous representations, the reward function is subjective: Cheese and shocks are only good and bad in someone's head. Thus, the specification of subjective rewards could also be manipulated by norms.

So MDPs describe decision situations with four variables: the set of states S , the actions at each decision point A , the transition function T , and the reward function R . In plain English, the person needs to know what the states of the world are, what she can do at each state, what next state she will arrive at, and how good or bad the consequences are. All four of these variables could, in principle, be influenced by norms in a way that would promote compliance.

RL Glossary

Reward: A subjective measure that the agent is attempting to maximize. In computer science this is defined by the programmer; in biological systems, it might be a consequence of biological or cultural evolution.

State: An agent's representation of the present circumstances relevant to decision-making.

Action: A behavioral response available to an agent in its current state.

Transition function: A function describing the probability of attaining a particular subsequent state given the current state and action.

Policy: The current likelihood with which the agent will decide to perform any given action in any given state—i.e., a complete probabilistic description of what the agent tends to do.

Value: The future reward that an agent expects to obtain, conditioned upon a given action and its current state and policy. I.e., "if I do X , and given the actions I am likely to take subsequently, how much reward can I expect to obtain?"

The goal of reinforcement learning is to accurately estimate the value of actions ("learning"), and then to choose actions that tend to have the highest value ("decision-making"). More formally, the agent wishes to converge on a reward-maximizing policy.

to current rewards (Sutton & Barto, 1998).

2.2 Discovering solutions: Reinforcement learning

Reinforcement learning is a branch of computer science that seeks algorithms for learning and deciding within MDPs. At the most abstract level, the optimal approach is for an agent to compute the *value* of each action available at each state, which has a precise definition: the expected sum of future rewards conditioned on choosing that action at the current decision point. Then, the person can choose the action with the highest value. This is analogous to computing the expected utility of a gamble in economic decision theory (von Neumann & Morgenstern, 1944).

Although value can be defined simply, calculating it is complicated. In a sequential setting such as the MDP, the long-term rewards following from an agent's present decision depend not only on that decision, but also on the *subsequent* decisions that the agent makes. In other words, the agent needs to know something about itself—for a chess algorithm to know which piece to move, it needs to consider what subsequent moves it is likely to implement as the game unfolds. This can be formalized via the concept of a "policy", which is a complete probabilistic mapping from state to action; a description of the likelihood of the agent doing each possible thing in each possible set of circumstances. Thus, for an agent to assign value to an action, it must condition not only on the present state but also on its own policy. From this vantage point, the agent's overarching goal is to develop an optimal policy.

A consequence of this complexity is that rational decision-making is often intractable. Consequently, much work in reinforcement learning seeks to find computationally efficient mechanisms for approximating value representation. Broadly speaking, many algorithms can be positioned on spectrum from those that are highly rational (achieving accuracy at the cost of computation) to those that are highly heuristic (sacrificing accuracy to lower computational demands). This distinction forms a major dividing line between rival theories of norm compliance, and so it worth our careful attention.

The standard economic approach defines the rational end of this spectrum: Use an explicit model of the situation to perform the exact calculations. For example, if Lenny knows that turning left twice leads him to cheese and turning right twice leads him to shock (i.e. he

has a representation of the transition function T), he can compute that the expected value of turning left at the initial fork is +1. In MDPs, this is known as goal-directed planning, and captures the kind of forward thinking often associated with human rationality (Elster, 1989; Harsanyi, 1977; Kahneman & Thaler, 2006). In the RL literature these methods are typically described as "model-based", because they require an agent to learn and explicitly represent a model of the transition structure of the environment, which forms the cognitive basis for planning. This approach can be very precise—even perfectly so—but it can also entail major computational costs.

Alternatively, Lenny could approximate the value of turning left at the fork by keeping a running average of the reward he's received after taking that action in the past. In this case, he wouldn't turn left because he believed it would lead to cheese; he would turn left simply because it had been a good choice in the past. In other words, Lenny would be acting habitually. This backward-looking method is often associated with irrationality (Elster, 1989)—or, more charitably, with heuristic cognition (Tversky & Kahneman, 1974)—and there is extensive evidence that people maintain and rely on these "cached" averages of past reward (Dolan & Dayan, 2013). In the RL literature, these methods are typically described as "model-free", because they do not require learning, representing or planning over a model of the transition structure of the environment; they just rely on averages of past reward. This approach often sacrifices precision but achieves major computational savings. We denote $Q(s,a)$ as the average value an agent has obtained from choosing action a in decision point s in the past. Q 's are the fifth and final decision variable in the RL framework that could be influenced by norms.⁴

The computational distinction between "model-based" and "model-free" approaches—i.e., between habit and planning—helped to

⁴ This method of habit learning by reinforcing actions historically associated with reward emerged in the reinforcement learning literature in the 1980s and rapidly revolutionized the field, quickly enabling human-level proficiency in games like backgammon. Several computational signatures of these model-free RL algorithms were also discovered in dopaminergic neural circuits that implement value-guided learning and decision-making, catalyzing two decades of rapid theoretical and empirical advances (Schultz, Dayan, & Montague, 1997).

crystalize a more fluid set of concepts that have long been fundamental to behavioral research. Psychologists since Thorndike (1927) have debated the extent to which complex behavior is the result of planning over an internal representation of the causal structure of the world, or instead the accumulation of adaptive heuristics shaped under the prior history of reward and punishment. RL offers one means of describing the competition between rival cognitive "systems" that are characterized by different tradeoffs of accuracy against computational demand (Kahneman, 2011).

The conflict between habit and planning is spelled out in a simple behavioral paradigm that mirrors our example of Lenny. A rat is trained to press a lever in order to obtain a food reward. Then, it is removed from the box and given free access to as much of its favorite foods as it wants, until it is completely sated and shows no further interest in the food. Finally, the sated rat is returned to the box with the lever. If the rat had received small amounts of training, it will not push the lever because it knows it doesn't want the food (i.e. it engages in planning). However, if it was overtrained and formed a habit of pressing the lever, it will move to the lever and continually press it, ignoring the food that accumulates on the floor (Dickinson, Balleine, Watt, Gonzalez, & Boakes, 1995). The irrational persistence in an action that was previously reinforced is the hallmark of habitual control of behavior. Reinforcement learning offers a means of making this distinction between habit and planning computationally precise, and assimilates a range of previously disparate decision making models into a common framework (Dolan & Dayan, 2013).

2.4 Summary

Together, MDPs and RL offer a powerful model of decision situations, with five variables (S , A , T , R , and Q) that could be influenced by norms (Fig. 2). This framework will form the backbone of our discussion. But it is important to emphasize again that this is not the only viable model of decisions. There are, for example, many models that forgo the notions of reward

and value, and don't assume all objectives can be mapped onto a common currency (Lichtenstein & Slovic, 2006; Vlaev et al., 2011). We adopt a reward-based framework because it is widespread in both cognitive decision theory and theories of norms, but future work might pursue these alternatives. There are also models of norm compliance that forgo any notion of mental representations at all (McClelland, Rumelhart, & Hinton, 1986). While non-representational models have much to contribute, our goal is to see how much traction cognitive models can gain.

3. A common language for theories of norm compliance

We use the framework of MDPs and RL algorithms to couch four major theories of norm compliance in a common language (Fig. 2).

3.1 Compliance by planning: A folk theory of norm influence

The simplest theory of norm influence was stated by Glaucon in Plato's *Republic*: People are "are only diverted into the path of justice by... force" (Jowett, 1925). Paraphrasing: People follow norms because they fear reprisal, or a loss of reputation. Although Glaucon's specific theory focused on punishment, we are concerned not with the *content* of people's expectations ("I'll go to jail"), but rather the *form* of the influence upon decision-making. The claim is that people obey norms because they explicitly represent the likely outcomes of compliance versus non-compliance, and conclude that the likely outcomes of compliance hold a greater prospect of maximizing long-term reward. Thus, this theory encompasses a wide range of specific motivations: Avoiding jail, earning respect, winning friends, maximizing profit, etc. This view is common enough to be called the folk theory of norm influence, and is the launching point for our discussion.

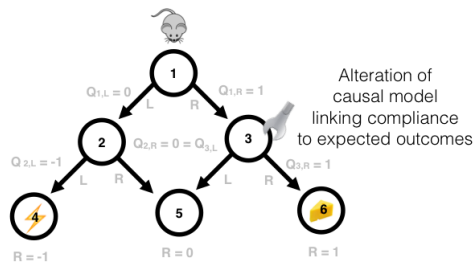
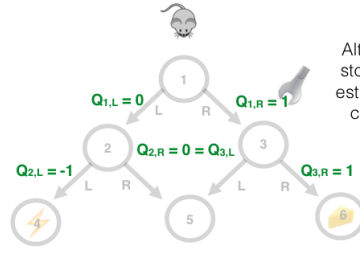
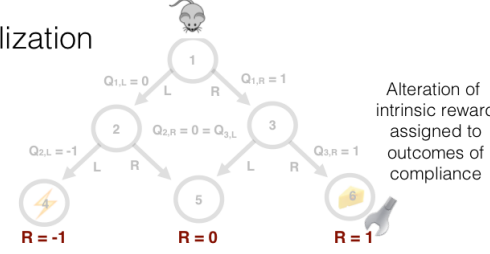
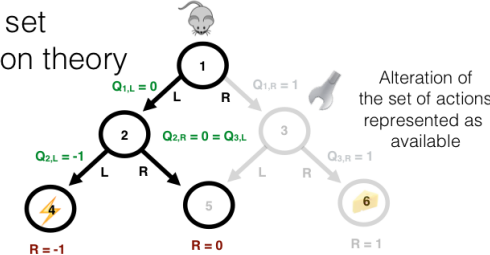
<p style="text-align: center;"><u>Variables</u></p>	<p style="text-align: center;"><u>Theories</u></p>
<p style="text-align: center;">Transition function $T(s,a,s')$</p>	<p style="text-align: center;">Folk Theory</p> 
<p style="text-align: center;">Cached value estimates $Q(s,a)$</p>	<p style="text-align: center;">Habit Theory</p> 
<p style="text-align: center;">Reward function $R(s,a,s')$</p>	<p style="text-align: center;">Internalization Theory</p> 
<p style="text-align: center;">Action set A_s</p>	<p style="text-align: center;">Choice set restriction theory</p> 
<p style="text-align: center;">State space S</p>	<p style="text-align: center;">No theory yet</p>

Figure 2: Five cognitive decision variables that norms could influence, and the theory of norm compliance that corresponds to each variable. On the folk model, norms change the decision maker’s internal causal model of the world (“if I cheat on my taxes, I’ll go to jail”). On the habit model, norms change the stored values of options via day-to-day reward and punishment of obedience or disobedience (“every time I cheated on things like taxes in the past, bad things happened”). On the internalization model, norms affect the intrinsic reward that people assign to different outcomes (“it’s wrong to not pay your fair share of taxes”). And on the unthinkable action model, norms change which actions are even considered (“I didn’t even think to cheat on my taxes”).

Construed within the RL framework, the folk theory comprises two key claims:

1. The agent's *expectation* that future reward will be maximized via norm compliance, by altering the agent's internal representation of the transition function T .
2. The existence of the norm does not change the agent's *assignment* of reward (i.e. the reward function R). In other words, norm compliance has value to the agent exclusively because it is linked to other rewards that the agent already holds (e.g., food, pain, etc.).

In the sections that follow, we will review alternative models of norm compliance that depart from the folk theory in two different ways. One claims that norms influence behavior not by planning over an explicit model of future reward, but rather by heuristics and habits established by the prior history of reward (i.e. norms affect Q). The other claims that norms change just not our expectation of reward but rather the very set of events that we find rewarding at all (i.e. norms affect R , e.g., by "internalization" of the norm as an intrinsic good). Thus, the folk theory stands apart from these alternatives principally in its claims of *planning* over an *unmodified* specification of reward.

It is uncontroversial that the folk theory captures at least part of the manner in which norms influence human behavior. At least some of the time, people pay taxes because they fear an audit; vote because they want to influence the outcome; tip because they want good service; recycle because they want to burnish their social image; etc. But it is equally uncontroversial that the folk theory is incomplete. The problem is that people routinely obey norms even when they know there is little chance of reprisal. They report their taxes honestly when the risk of an audit is incredibly low; vote in elections which they have no chance of influencing; tip in restaurants they will never revisit; recycle, even when living alone; and so on (Sunstein, 1996).

The persistence of norm compliance in the absence of explicit expectations of enforcement has been most carefully demonstrated with laboratory economic games (Camerer, 2003). In these games, people are given an endowment of bonus money and asked to make decisions that will affect both their own payment and the payment of other subjects. For example, people

are sometimes asked whether they want to contribute to a public pool that will benefit everyone else but cost them personally. Other times they are asked to split the endowment with another subject. Time after time, most participants at least sometimes obey norms prescribing fairness, cooperation, reciprocity, or honesty. Importantly, they follow the norms even when the game is one-shot and anonymous—in other words, when their actions are guaranteed to go unpunished (Fehr & Schmidt, 2006).

We call this phenomenon—people's dogged persistence in obeying norms—the *persistence problem*. People don't always persist in norm compliance, but they do it often enough to make us question whether explicit expectations of reward are truly the driving force. The folk theory can be thought of as a null hypothesis, and the persistence problem the data which force us to consider the more complex alternatives that follow.

There are a few concerns with this analysis that will be useful to address before moving to more complex theories. First, is the persistence problem really a problem, or just an ordinary mistake? People could simply be failing to notice that they are in a situation where punishment doesn't apply, or fail to incorporate this knowledge into their representation of the transition function (Delton, Krasnow, Cosmides, & Tooby, 2011). Though this objection is difficult to rule out in real-world examples, the simplicity and transparency of countless laboratory experiments makes it unlikely that subjects are not representing the contextual one-shot anonymity (Andreoni, 1995; Fehr & Schmidt, 2006; Thaler, 1988). Moreover, people *do* often incorporate this information into their planning, increasing compliance when people are observing them or can reciprocate (Fehr & Schmidt, 2006). The interesting fact is that their levels of norm compliance are far from zero even in the one-shot anonymous case.

Second, we have emphasized examples like tax reporting, recycling, and one-shot economic games to draw inferences about norm compliance. But how do we know that these are actually examples of norm compliance? People might have performed these behaviors in the absence of exposure to any norm (e.g. because of evolved instinct). Again, while this is difficult to rule out in the real world, numerous experiments show that norms drive people's behavior in anonymous situations. People

exhibit enormous cross-cultural variation in one-shot economic games, and this variation is partially explained by differences in cultural norms (Henrich, 2000; Henrich et al., 2006, 2010). And manipulating norms changes behavior. People walking alone rarely litter after seeing anti-littering signs, but litter quite often after seeing piles of trash on the ground (Cialdini, Kallgren, & Reno, 1991; Cialdini, Reno, & Kallgren, 1990). People reuse hotel towels more often, and give away more money in economic games, when told it is the norm (Goldstein, Cialdini, & Griskevicius, 2008; Raihani & McAuliffe, 2014). All this suggests that people are influenced by norms, even when they won't be sanctioned for disobedience—and thus the folk theory of norm compliance is incomplete.

To motivate our exploration of the theory space, we will repeatedly draw on these kinds of examples as genuine instances of norm-following. The key assumption in each instance is that the aggregate behavior—a tendency to offer fair splits of a monetary endowment, contribute to a public good, return an investment, etc.—is in part driven by prior exposure to norms. Thus, from patterns in the behavior, we can infer features of how norms influence us. In many cases, there is experimental (and anecdotal) evidence to support this assumption. But in other cases, it has not been explicitly tested.³ We will highlight claims that demand more rigorous experimental support from future work.

Finally, it is important to note that our review of literature will be dramatically incomplete. We focus on studies that we believe offer the most precision and insight into the cognitive underpinnings of compliance, and, in doing so, pass over much important work in sociology and other fields. We hope to offer a cognitive perspective that naturally complements these alternative approaches.

³ Many studies show that manipulating a variable (e.g. time available to make a decision) makes people more or less cooperative, fair, generous, etc., in economic games. These choices are likely influenced by the norms to which people have been exposed (Fehr & Fischbacher, 2004; Rand et al., 2014). But ideally, to show that the specific variable at hand affects norm compliance, the study would simultaneously manipulate whether the relevant norm is present, and show that the variable only has an effect when the norm is present. Unfortunately, the ideal study has often not yet been run. We cite the imperfect studies with the hope that future work will fill in the gaps.

3.2 Compliance by habit: Heuristic approaches to value approximation

The folk theory assumes that people are representing their environment's transition function ("if I turn right at the fork, it will lead to cheese"), and selecting actions via forward planning ("I want cheese, therefore I should turn right"). They know the consequences of their choices, and act on that knowledge. This assumption is what causes trouble for the folk theory: People seem to follow norms even when they know the typical motivating consequences are suspended (e.g. in one-shot anonymous settings).

There are two ways out of this bind. Perhaps the consequences that motivate norm compliance are not sanctions or reputation, but an intrinsic "social preference"—e.g. promoting fairness or being a good person. This solution to the persistence problem depends on altering the reward function R , and it is pursued in the next section.

First, though, there is a simpler alternative. As described above, people don't always plan ahead; they often act out of habit. More formally, instead of representing the transition function T (i.e. the explicit consequences of actions), they store the average value Q of past actions, and simply choose actions with higher Q s—actions that accumulated higher reward in the past. A habitual Lenny turns left in the maze, not because he knows it will lead to cheese, but because he remembers it led to an average goodness of +1 in the past. If you told him that the cheese had moved, it wouldn't affect his immediate decision making, because he's not even considering cheese when making his decision.

Similarly, a habitual human might propose a fair monetary split in an economic game, not because she knows it will avoid sanctions and improve her reputation, but because she merely remembers that being fair led to an "average goodness" of +1 in the past. If you told her that the game was one-shot/anonymous and she couldn't be sanctioned, it wouldn't affect her immediate decision making, because she's not even considering sanctions when making her decision.

We call this the habit theory of norm compliance. In everyday life, most situations are not one-shot or anonymous, and complying with norms is typically beneficial (to avoid

sanctions, accrue reputation, and so on). This experience produces high cached Q values for norm-compliant actions. Then, when people are in one-shot anonymous settings where the typical benefits are removed, they choose based on those Q values, which are not immediately affected by the knowledge of anonymity. Thus, people persist in norm compliance (Rand et al., 2014; Rand, Greene, & Nowak, 2012).

One point in favor of the habit theory is that people appear habitually averse to causing direct personal harm, even in contexts where it maximizes welfare (Crockett, 2013; Cushman, 2013). In contemporary societies, there are strong norms against causing typical forms of harm: pushing, hitting, shooting, stabbing, etc. The habit theory predicts that people will therefore have developed a habitized aversion – i.e. a low Q value – to performing those behaviors. In support of this, people are unwilling to perform these actions (e.g. pushing someone in front of a runaway trolley), even in the presence of a causal model linking such an action to positive consequences (e.g. saving five other people; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001). Moreover, people exhibit physiological signs of aversion to pretend actions that resemble acts of physical harm, such as pulling the trigger of a fake replica handgun while pointing it at a person's head (Cushman, Gray, Gaffey, & Mendes, 2012). Like the rat who keeps pushing its food lever even when full, people seem to be relying on a cached value that they have assigned to certain norm-violating behaviors – even when the typical consequences of those behaviors are suspended (Cushman, 2015).

The habit theory of norm compliance is also supported by a suite of experimental results motivated by the fact that acting on habits requires less computation time/resources than planning. (Planning requires the agent to integrate over a potentially complex transition function, while acting on habits just requires a comparison among pre-computed Q values.) Thus, according to the habit theory, people with less time or mental resources should be more compliant.

Indeed, when people in economic games are put under time pressure or cognitive load, they become more cooperative (Rand, 2016), fair (Cappelletti, Güth, & Ploner, 2011; Halali, Bereby-Meyer, & Ockenfels, 2013), reciprocal (C. Anderson & Dickinson, 2010; Grimm & Mengel, 2011; Halali, Bereby-Meyer, & Meiran, 2011,

2014; Neo, Yu, Weber, & Gonzalez, 2013; Sutter, Kocher, & Straub, 2003), and—sometimes—more generous (Schulz, Fischbacher, Thöni, & Utikal, 2014).⁶ This finding resonates with real-world reports of extreme prosociality: People who risk their lives to help others almost always describe their decision process as quick and intuitive (like acting on a habit), not slow and reflective (like planning; Rand & Epstein, 2014).

Moreover, the effect of time pressure / cognitive load in economic games is dependent on prior experience: It is strongest for people who report positive cooperative experiences in their daily life (Rand et al., 2014, 2012), and disappears with extensive experience in one-shot anonymous economic games (Rand et al., 2014). Both of these findings are explained by the habit theory. Only people who were rewarded for complying with norms in the past should form a habit of it. And, while Q values are not *immediately* affected by contextual anonymity, they will be in the long run: Repeated experience in a context where compliance is not beneficial will drive down the Q values until compliance is no longer a habit. Imagine we move Lenny's cheese. If Lenny keeps habitually turning left in the maze, eventually the average reward for turning left will plummet, and Lenny will stop turning left. Similarly, if people keep playing one-shot anonymous economic games, their average stored reward for complying with norms will drop, they will stop complying out of habit, and thus the effect of time pressure will disappear. These experimental results therefore support the habit theory of norm compliance.⁷

⁶ The effect of time pressure and cognitive depletion on generosity is mixed. Some studies report that they induce more giving (Cornelissen, Dewitte, & Warlop, 2011; Schulz, Fischbacher, Thöni, & Utikal, 2014), while others report a null or (rarely) reversed effect (Hauge, Brekke, Johansson, Johansson-Stenman, & Svedsäter, 2015; Halali, Bereby-Meyer, & Ockenfels, 2013). Interestingly, the game used in these studies to measure generosity has notoriously fickle norms (Bicchieri, 2006; Fehr & Schmidt, 2006; Krupka & Weber, 2013). Perhaps the ambiguous effects of time pressure on generosity can be explained by differences in norm perception across studies.

⁷ It is possible that, under time pressure, people do not become more compliant; they simply become more prosocial (Rand, Greene, & Nowak, 2012). The fact that people are also more negatively reciprocal under time pressure suggests the former interpretation. But it is an open question. One study appeared to show that, even after being instilled with a norm for competition instead of cooperation, people were still more cooperative under time pressure (Cone & Rand, 2014). But people in the competitive condition did

In sum, habit theory likely explains some part of norm compliance. But it is also incomplete. The problem is that, even after people have enough experience with one-shot anonymous settings to lose their habit for compliance, they continue to comply. In the real world, experienced travelers still tip in foreign cities; experienced solo-hikers still avoid littering when alone in public parks; hotel guests still reuse towels on the tenth day of their stay. In all these cases, if people were acting out of a habit developed in non-anonymous settings, experience with the anonymous setting should eliminate compliance—but it doesn't.

This phenomenon is most precisely demonstrated with laboratory games. As described above, people who participate in many economic games seem to lose their habit for compliance; they stop showing an effect of time pressure on compliance choices. Yet they continue to contribute to public goods, propose fair monetary splits, return investments, etc., at almost the same rate (Rand, 2016; Rand et al., 2012)—it just no longer matters whether they have a short or a long amount of time to do so. The persistence of norm-compliant behavior even in the face of massive experience with one-shot, anonymous contexts is difficult to explain under the habit theory.

In sum, the habit theory appears only to explain "shallow" persistence: People comply with norms in the first few anonymous situations they encounter. But it does not naturally account for "deep" persistence: People continue to comply with norms, even after they have enough experience with anonymous situations to kick the habit. Norms seem to do more than just instill habits.

3.3 Compliance by internalization: Norms as sources of intrinsic reward

A popular idea, hinted at before, is that norm compliance may sometimes be valued intrinsically, rather than as an instrumental means of avoiding punishment or attaining reward. Through socialization, people internalize the norm and come to care about the things that norms prescribe. People are fair, not because they know unfairness would be punished, but because they care about

producing fair outcomes (Fehr & Schmidt, 1999). They contribute to public goods, not because their contributions are being externally tracked, but because they believe it's the right thing to do (Cialdini & Trost, 1998; Parsons & Shils, 2011). If this were true, then of course people would continue to comply with norms in anonymous settings—anonymousity doesn't remove the consequences they care about.

This concept can be made precise within the reinforcement learning framework by positing that that exposure to norms changes people's reward functions, R (Ho, MacGlashan, Littman, & Cushman, 2017). Recall that the function of reinforcement learning is to estimate the instrumental *value* of an action based on the long-term expectation of intrinsic *reward* in subsequent states. Whereas the folk theory and habit theory posit that norms affect behavior by altering the way that we perceive the instrumental value of compliance, the internalization theory posits that norms affect behavior by altering the very set of states or actions that we find intrinsically rewarding.

This theory naturally accounts for deep persistence without requiring us to abandon the possibility of rational planning (Fehr & Fischbacher, 2003). In the language of MDPs, the dictates of the norm ("split goods fairly") have just been internalized as another part of people's reward functions (e.g. +1 for proposing a fair split in an economic game), which can be flexibly and rationally traded off against other parts of the reward function during planning. Internalization theory also explains the role of guilt (and other negative emotions) in preventing norm violations: Guilt is part of the internalized negative reward for violating norms (Fehr & Schmidt, 2006).

Although the internalization theory can be stated in terms of MDPs, it is important to note that it falls outside the scope of the RL algorithms typically pursued in machine learning contexts. Ordinarily, a reinforcement learning agent operates with an unchangeable reward function defined by the programmer. It should be obvious why: If the programmer specifies that winning chess is the reward for an agent, it would not want that agent to later redefine rewards around, say, making pretty patterns with the pawns. Yet this, more or less, is what internalization theory posits: That humans initially designed to maximize one thing (e.g., maximum food for themselves) might later autonomously decide to try to

not, on average, contribute less than people in the cooperative condition (the difference was around 1 cent out of an endowment of 40), suggesting that the norm manipulation was insufficient.

maximize something else (e.g., equal food for all). An interesting area of contemporary research seeks to understand why natural selection might favor such an architecture, in which reward itself is subject to evaluation and change (Ackley & Littman, 1992; Ho et al., 2017; Singh, Lewis, & Barto, 2009).

Although the internalization theory does some violence to traditional assumptions of reinforcement learning, its commitments can also be made much more precise within the RL framework. For instance, there is disagreement in the literature over *what*, precisely, gets internalized during norm learning. Some theories posit that humans learn to intrinsically value certain states of affairs, such as an equitable division of resources among people (Fehr & Schmidt, 1999). Other theories posit that humans learn to intrinsically value certain actions, such as the act of generating an equitable division (Andreoni, 1990). These have been regarded as competitors, yet often the precise commitments of various theories are not transparent.

A theory of norm internalization couched in the language of MDPs is naturally committed to precise answers to all of these questions, simply by specifying a reward function. The reward function R takes as input the current state s , the chosen action a , and the subsequent state s' . If norm internalization made the outcome of compliance rewarding (us receiving an equal amount of money), then R would be positive for an s' with equal splits, for any a (Fehr & Schmidt, 1999). If it made the act rewarding, then R would be positive for an a of proposing an equal split, for any s' (Andreoni, 1990). If the reward was for specific norms, then R would be defined over a set of concrete actions a ; if the reward was for following norms in the abstract, then R would be defined over an abstract action space of "follow the norm" or "don't follow the norm" (Dana, Cain, & Dawes, 2006; Krupka & Weber, 2013; López-Pérez, 2008). If context mattered, then R would be positive for only certain states s . MDPs can precisely model all of these situations, showcasing the breadth of ideas that fall under the umbrella of the reward theory of compliance.

3.5 A challenge to value-guided theories

Each of the theories that we have considered so far shares a common assumption: Norms intervene on decision-making by altering the value of actions under consideration, as they

compete against alternatives in a common currency. This would appear to be basic commitment of any theory of norm compliance grounded in value-guided decision-making.

An important challenge to this perspective is that people sometimes disregard, or are highly insensitive to, the costs and benefits of complying with or violating these norms. For example, many people say there is no amount of benefit that would justify selling organs, auctioning off brides, or mandating abortions to limit population growth (Baron & Spranca, 1997). If a hospital director is deciding whether to perform a surgery to save a child's life, he should not consider the surgery's cost (Tetlock, Kristel, Beth, Green, & Lerner, 2000); if a soldier is deciding whether to obey orders, she should not consider the benefits of disobedience (Raz, 1999). These thought experiments about "taboo tradeoffs" or "incommensurate values" are so compelling that some philosophers consider the inability to trade off compliance with other goods to be the essence of moral norms (Nozick, 1974; Raz, 1999; Schauer, 1991). This suggests that the motivation to comply with moral norms is not always of a common currency with other goods, and is not simply incorporated into people's reward functions.

Two objections to this line of argument immediately come to mind. First, perhaps people place arbitrarily large negative reward on the outcome of certain norm violations, accounting for the apparent "unthinkable" nature of non-compliance within a value-guided framework. For example, they might place infinitely negative reward on a state in which abortions are mandated, or organs are sold on markets. Thus, any benefit of violating the norm would be outweighed by the cost, and people will be unwilling to violate it. This interpretation appears to reconcile protected values with the reward theory of norm compliance, obviating the need for an alternative theory.

The consilience, however, is misleading. Rewards are common-currency, numerical representations of the subjective, intrinsic goodness in certain states or actions. An essential feature of such representations is that they can be traded off with each other (Becker, 1996). Saying that a reward is too large to ever be outweighed is equivalent to saying that it is not a reward.

Second, even though people *say* they are unwilling to trade off compliance with other

goods, perhaps they still would if the right situation arose. Perhaps their talk of taboos is just signaling ("I would *never* do that!"), or they are not imaginative enough to consider situations with truly enormous benefits of disobedience (Tetlock et al., 2000). After all, when prompted, juries are able to put a price on something as sacred as a human life (Ubel, 2001). More work needs to be done to test whether, in their actual decisions, people are unwilling to put norm compliance on a common scale with other goods.

There is reason, however, to think that people genuinely treat compliance differently from other goods. People often go to absurd lengths to comply with trivial situational norms, far beyond what a reasonable cost-benefit, common-currency analysis would recommend. For instance, subjects in the Milgram experiment chose to fatally electrocute another human rather than disobey norms of proper conduct (Milgram, 1963). Do people value norms of proper conduct more than human life? More likely, the norms were influencing behavior through some non-value-guided channel, thus circumventing this comparison. What is this alternative channel?

3.6 Compliance by choice set construction: A model of "unthinkable" action

Here, we suggest one possibility. The basic insight of the "taboo tradeoffs" literature is that people sometimes take norm violation off the table completely as an option. When people are unwilling to consider the benefits of organ markets, they are effectively eliminating selling their kidney from the set of choices they consider. On this view, moral norms aren't just another factor in the expected value of an option; they influence which options even get their expected value calculated in the first place. They change which options get entertained (Fig. 2).

This theory can be naturally formulated in the language of MDPs. A crucial variable in an MDP is the set of available actions A at each decision point s . This variable is subjective, and must be constructed by the decision maker. It corresponds to what economists sometimes call the "choice set" that a decision-maker constructs prior to choice, narrowing from all conceivable actions to a restricted set of considered alternatives. The idea is that norms influence choices in part by excluding norm-violating actions (e.g. selling your kidney) from A , for

some broad set of decision points s . Thus, after being exposed to a norm, people become less likely to violate it.

As before, formulating the theory in the language of MDPs allows us to clarify some of its ambiguities. It would be foolish to claim that moral norms always eliminate norm-violating actions from A —people sometimes do sell their kidneys, and soldiers sometimes do disobey orders. The theory needs some kind of hedge to account for these exceptions. But where exactly should it hedge? One possibility is that the process that generates the action set A is probabilistic, and norms make it less likely (but not impossible) for norm-violating actions to make the cut. Alternatively, being exposed to a norm could cause you to deterministically exclude certain actions from A , but only in certain states s (e.g. never consider selling your kidney, unless you are in a state of desperation). MDPs can capture both possibilities.

There is also ambiguity in what it means to exclude an action from A . It could mean that the decision maker is aware of the action, but deems it irrelevant. This view is bolstered by the fact that people often deem norm-violating options irrelevant when judging others. For example, consider a captain on a sinking ship who throws either some cargo, or his wife, off the ship. In the latter case, people think the captain acted freely, because he could have thrown the cargo off instead. But in the former case, people think the captain was forced, because there was nothing else he could have done (Phillips & Knobe, 2009). This belief persists even if people are reminded that he could have thrown his wife off the ship instead; they say that option is irrelevant. This tendency to deem norm-violating options irrelevant seems to appear whenever people judge others' choices (Phillips, Luguri, & Knobe, 2015). Perhaps people apply the same tendency to their own choices, and deem norm-violating options irrelevant to the decision-making process.

Alternatively, excluding an action from A could mean that the decision maker is never aware of the action—it doesn't even come to mind. (More precisely, the representation of a norm-violating action ["I should sell my kidney..."] is never tokened in the person's mind.) While this seems unlikely in laboratory experiments with explicit, constrained choice sets (e.g. "cooperate" or "defect"), real-world decisions typically have choice sets that are unbounded and unstructured (think of all the things you could

conceivably do on a given Saturday afternoon). People never represent the vast majority of the choices they could make, and norms might influence which ones make the cut. The cognitive mechanisms behind choice set construction are extremely unclear, and so it is unclear how norms would produce this effect. But the idea is ripe for future research.

Why would a decision maker ever exclude actions from consideration? Why not consider all of them? In laboratory experiments with constrained choice sets, it seems implausible that people wouldn't evaluate every option. But again, in the real world, choice sets are enormous. People must have mechanisms for narrowing down their options to a small subset (Cushman & Morris, 2015; Huys et al., 2012). Since norm-violating options are typically unwise, people would obtain a practical benefit from excluding them.

Moreover, there is a strategic benefit to not even think about violating norms. If other people notice that you considered violating a norm (e.g. you took two minutes before deciding to cooperate in a repeated economic game), they'll infer that you *would* violate the norm if the incentives were strong enough, and thus terminate productive relationships with you. In other words, in repeated games where your partners know which options you consider, it is strategically beneficial to obey norms "without looking" (Hoffman, Yoeli, & Nowak, 2015; Jordan, Hoffman, Nowak, & Rand, 2016). These considerations further motivate the theory that norms cause people to ignore norm-violating options.

The possibility that norms exclude options from consideration stands out as an important new frontier for research. Some preliminary evidence comes from a study showing that, by default, people tend to treat immoral actions as if they were actually impossible (Phillips & Cushman, in press). Only after a little time and thought do they acknowledge that immoral actions are possible, albeit wrong. Although this pattern of response becomes most pronounced in adults subject to time pressure, it is promiscuous in the judgments of young children. For instance, at four years old, a majority of children claim that it is impossible to steal a candy bar from a store—and, moreover, that to succeed would require magic to do so (Phillips & Bloom, under review).

4. Conclusion

We identified four viable theories of how norms influence human behavior (Fig. 2). Each theory specified a representation involved in decision making that norms could alter. Norms could change the representation of the transition function T ("if I cheat on my taxes, then I'll go to jail"); the average values of past actions Q ("I've learned to associate cheating on things like taxes with badness"); the reward function R ("I would feel guilty for not paying my fair share"); or the action set A ("I never even considered cheating on my taxes").

Construing theories of norm compliance within a common framework has several benefits. We briefly highlight three of these. First, by formalizing norm compliance within the MDP setting, we can identify new potential mechanisms—ones that are currently poorly represented in the literature. For instance, one of these is the definition of the state space, S . Consider, for instance, the norm of impartial beneficence—i.e., the claim that welfare benefits should be distributed identically to all people without respect to their personal identity or group membership. This norm might naturally be captured by dictating that the state representations implicated in welfare decisions should also be insensitive to personal identity or group membership. This formalizes the ordinary notion of "blind justice", in the sense that rules are applied in a manner deliberately isolated from a source of potential information. Despite its intuitive appeal, we know of little current research on the interface between norm learning and state representation in the psychology, neuroscience or economics literatures.

Second, drawing on the MDP formalization facilitates comparison between norm compliance and other forms of decision-making. The MDP formalization and several RL algorithms defined within it have enabled notable advances in psychology and neuroscience. This research has mostly focused on the principles of learning in non-social environments. There is considerable interest in determining whether the influence of morality—and, more broadly, social norms—operates via overlapping or even identical mechanisms (Cushman, 2013; Ruff & Fehr, 2014). Asking this question requires us to map theories of social and non-social decision-making into a common conceptual space, as we have attempted here.

Third, the framework draws out certain unexplored points of divergence between the rival theories. For instance, by positing that norms affect different decision variables, the theories implicitly make different claims about how easy it is to change people's norm-compliance tendencies. The transition function T is relatively pliable; it flexibly incorporates many sources of information (like verbal instructions; Daw, Gershman, Seymour, Dayan, & Dolan, 2011). Thus, according to the folk theory, it should be easy to change people's norm-compliance tendencies. In contrast, the cached values Q are thought to primarily reflect direct experience with rewards and punishments; and the reward function R is often assumed to be nearly immutable. Thus, if norms affect Q (habit theory) or R (internalization theory), their influence is less easily undone. This highlights just one of several ways in which the MDP formalization can draw out new insights about the nature of rival theories of norm compliance.*

As should be apparent, by exploring the space of possible cognitive mechanisms underlying norm compliance, we have raised more questions than we answered. Nonetheless, we hope that this review will encourage the kind of cross-pollination necessary to progress our understanding of this vital topic. Norms touch every aspect of our external lives; one day, we may know which parts of our internal lives they touch as well.

Acknowledgements

We thank Jonathan Phillips and other members of the Moral Psychology Research Lab for their advice and assistance. This research was supported by Grant N00014-14-1-0800 from the Office of Naval Research.

References

- Ackley, D., & Littman, M. (1992). Interactions between learning and evolution. *Artificial Life, II*.
- Anderson, C., & Dickinson, D. L. (2010). Bargaining and trust: the effects of 36-h total sleep deprivation on socially interactive decisions. *Journal of Sleep Research, 19*(1-Part-I), 54–63. <https://doi.org/10.1111/j.1365-2869.2009.00767.x>
- Anderson, C. M., & Putterman, L. (2006). Do non-strategic sanctions obey the law of demand? The demand for punishment in the voluntary contribution mechanism. *Games and Economic Behavior, 54*(1), 1–24. <https://doi.org/10.1016/j.geb.2004.08.007>
- Andreoni, J. (1990). Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving. *The Economic Journal, 100*(401), 464–477. <https://doi.org/10.2307/2234133>
- Andreoni, J. (1995). Cooperation in Public-Goods Experiments: Kindness or Confusion? *The American Economic Review, 85*(4), 891–904.
- Andreoni, J., Castillo, M., & Petrie, R. (2003). What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game. *The American Economic Review, 93*(3), 672–685.
- Andreoni, J., & Miller, J. (2002). Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism. *Econometrica, 70*(2), 737–753. <https://doi.org/10.1111/1468-0262.00302>
- Andreoni, J., & Vesterlund, L. (2001). Which is the Fair Sex? Gender Differences in Altruism. *The Quarterly Journal of Economics, 116*(1), 293–312.
- Baron, J., & Spranca, M. (1997). Protected Values. *Organizational Behavior and Human Decision Processes, 70*(1), 1–16. <https://doi.org/10.1006/obhd.1997.2690>
- Becker, G. S. (1996). *Accounting for Tastes*. Harvard University Press.
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Camerer, C. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.
- Cappelletti, D., Güth, W., & Ploner, M. (2011). Being of two minds: Ultimatum offers under cognitive constraints. *Journal of Economic Psychology, 32*(6), 940–950. <https://doi.org/10.1016/j.joep.2011.08.001>
- Capraro, V., Jordan, J. J., & Rand, D. G. (2014). *Heuristics Guide the Implementation of Social Preferences in One-Shot Prisoner's Dilemma Experiments* (SSRN Scholarly Paper No. ID 2429862). Rochester, NY: Social Science Research Network. Retrieved from <http://papers.ssrn.com/abstract=2429862>
- Carpenter, J. P. (2007). The demand for punishment. *Journal of Economic Behavior &*

* An alternative formulation of this idea is that the decision variables have different levels of informational encapsulation (Fodor, 1983): Q and R are more encapsulated than T , and there are therefore fewer types of experience that can change them. At the present, it is unknown how encapsulated the action set A or the state space S are.

- Organization*, 62(4), 522–542.
<https://doi.org/10.1016/j.jebo.2005.05.004>
- Cialdini, R. B., Kallgren, C. A., & Reno, R. R. (1991). A Focus Theory of Normative Conduct: A Theoretical Refinement and Reevaluation of the Role of Norms in Human Behavior. In M. P. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 24, pp. 201–234). Academic Press. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0065260108603305>
- Cialdini, R. B., Reno, R. R., & Kallgren, C. A. (1990). A focus theory of normative conduct: Recycling the concept of norms to reduce littering in public places. *Journal of Personality and Social Psychology*, 58(6), 1015–1026.
<https://doi.org/10.1037/0022-3514.58.6.1015>
- Cialdini, R. B., & Trost, M. R. (1998). Social influence: Social norms, conformity and compliance. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *The handbook of social psychology, Vols. 1 and 2 (4th ed.)* (pp. 151–192). New York, NY, US: McGraw-Hill.
- Cone, J., & Rand, D. G. (2014). Time Pressure Increases Cooperation in Competitively Framed Social Dilemmas. *PLOS ONE*, 9(12), e115756.
<https://doi.org/10.1371/journal.pone.0115756>
- Cooter, R. (2000). Do Good Laws Make Good Citizens? An Economic Analysis of Internalized Norms. *Virginia Law Review*, 86(8), 1577–1601.
<https://doi.org/10.2307/1073825>
- Cornelissen, G., Dewitte, S., & Warlop, L. (2011). Are Social Value Orientations Expressed Automatically? Decision Making in the Dictator Game. *Personality and Social Psychology Bulletin*, 0146167211405996.
<https://doi.org/10.1177/0146167211405996>
- Crockett, M. J. (2013). Models of morality. *Trends in Cognitive Sciences*, 17(8), 363–366.
<https://doi.org/10.1016/j.tics.2013.06.005>
- Cushman, F. (2013). Action, Outcome, and Value: A Dual-System Framework for Morality. *Personality and Social Psychology Review*, 17(3), 273–292.
<https://doi.org/10.1177/1088868313495594>
- Cushman, F. (2015). From moral concern to moral constraint. *Current Opinion in Behavioral Sciences*, 3, 58–62.
<https://doi.org/10.1016/j.cobeha.2015.01.006>
- Cushman, F., Gray, K., Gaffey, A., & Mendes, W. B. (2012). Simulating murder: The aversion to harmful action. *Emotion*, 12(1), 2–7.
<https://doi.org/10.1037/a0025071>
- Cushman, F., & Morris, A. (2015). Habitual control of goal selection in humans. *Proceedings of the National Academy of Sciences*, 112(45), 13817–13822.
<https://doi.org/10.1073/pnas.1506367112>
- Dana, J., Cain, D. M., & Dawes, R. M. (2006). What you don't know won't hurt me: Costly (but quiet) exit in dictator games. *Organizational Behavior and Human Decision Processes*, 100(2), 193–201.
<https://doi.org/10.1016/j.obhdp.2005.10.001>
- 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), 1204–1215.
<https://doi.org/10.1016/j.neuron.2011.02.027>
- Delton, A. W., Krasnow, M. M., Cosmides, L., & Tooby, J. (2011). Evolution of direct reciprocity under uncertainty can explain human generosity in one-shot encounters. *Proceedings of the National Academy of Sciences*, 108(32), 13335–13340.
<https://doi.org/10.1073/pnas.1102131108>
- Dickinson, A., Balleine, B., Watt, A., Gonzalez, F., & Boakes, R. A. (1995). Motivational control after extended instrumental training. *Animal Learning & Behavior*, 23(2), 197–206.
<https://doi.org/10.3758/BF03199935>
- Dolan, R. J., & Dayan, P. (2013). Goals and Habits in the Brain. *Neuron*, 80(2), 312–325.
<https://doi.org/10.1016/j.neuron.2013.09.007>
- Ellickson, R. (1994). *Order without Law: How Neighbors Settle Disputes* (Revised edition). Cambridge, Mass.: Harvard University Press.
- Elster, J. (1989). Social Norms and Economic Theory. *The Journal of Economic Perspectives*, 3(4), 99–117.
- Fehr, E., & Fischbacher, U. (2003). The nature of human altruism. *Nature*, 425(6960), 785–791.
<https://doi.org/10.1038/nature02043>
- Fehr, E., & Fischbacher, U. (2004). Social norms and human cooperation. *Trends in Cognitive Sciences*, 8(4), 185–190.
<https://doi.org/10.1016/j.tics.2004.02.007>
- Fehr, E., & Schmidt, K. M. (1999). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics*, 114(3), 817–868.
- Fehr, E., & Schmidt, K. M. (2006). Chapter 8 The Economics of Fairness, Reciprocity and Altruism – Experimental Evidence and New

- Theories. In S.-C. K. and J. M. Ythier (Ed.), *Handbook of the Economics of Giving, Altruism and Reciprocity* (Vol. 1, pp. 615–691). Elsevier. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1574071406010086>
- Fodor, J. A. (1983). *The Modularity of Mind: An Essay on Faculty Psychology*. MIT Press.
- Gershman, S. J., & Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*, 20(2), 251–256. <https://doi.org/10.1016/j.conb.2010.02.008>
- Glimcher, P. W., & Fehr, E. (2013). *Neuroeconomics: Decision Making and the Brain*. Academic Press.
- Goldstein, N. J., Cialdini, R. B., & Griskevicius, V. (2008). A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research*, 35(3), 472–482. <https://doi.org/10.1086/586910>
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI Investigation of Emotional Engagement in Moral Judgment. *Science*, 293(5537), 2105–2108. <https://doi.org/10.1126/science.1062872>
- Grimm, V., & Mengel, F. (2011). Let me sleep on it: Delay reduces rejection rates in ultimatum games. *Economics Letters*, 111(2), 113–115. <https://doi.org/10.1016/j.econlet.2011.01.025>
- Halali, E., Bereby-Meyer, Y., & Meiran, N. (2011). *When Rationality and Fairness Conflict: The Role of Cognitive-Control in the Ultimatum Game* (SSRN Scholarly Paper No. ID 1868852). Rochester, NY: Social Science Research Network. Retrieved from <http://papers.ssrn.com/ezip-prod1.hul.harvard.edu/abstract=1868852>
- Halali, E., Bereby-Meyer, Y., & Meiran, N. (2014). Between self-interest and reciprocity: The social bright side of self-control failure. *Journal of Experimental Psychology: General*, 143(2), 745–754. <https://doi.org/10.1037/a0033824>
- Halali, E., Bereby-Meyer, Y., & Ockenfels, A. (2013). Is it all about the self? The effect of self-control depletion on ultimatum game proposers. *Frontiers in Human Neuroscience*, 7. <https://doi.org/10.3389/fnhum.2013.00240>
- Harsanyi, J. C. (1977). Morality and the Theory of Rational Behavior. *Social Research*, 44(4), 623–656.
- Hauge, K. E., Brekke, K. A., Johansson, L.-O., Johansson-Stenman, O., & Svedsäter, H. (2015). Keeping others in our mind or in our heart? Distribution games under cognitive load. *Experimental Economics*, 19(3), 562–576. <https://doi.org/10.1007/s10683-015-9454-z>
- Henrich, J. (2000). Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining among the Machiguenga of the Peruvian Amazon. *The American Economic Review*, 90(4), 973–979.
- Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., ... Ziker, J. (2010). Markets, Religion, Community Size, and the Evolution of Fairness and Punishment. *Science*, 327(5972), 1480–1484. <https://doi.org/10.1126/science.1182238>
- Henrich, J., McElreath, R., Barr, A., Ensminger, J., Barrett, C., Bolyanatz, A., ... Ziker, J. (2006). Costly Punishment Across Human Societies. *Science*, 312(5781), 1767–1770. <https://doi.org/10.1126/science.1127333>
- Hoffman, M., Yoeli, E., & Nowak, M. A. (2015). Cooperate without looking: Why we care what people think and not just what they do. *Proceedings of the National Academy of Sciences*, 112(6), 1727–1732. <https://doi.org/10.1073/pnas.1417904112>
- Ho, M. K., MacGlashan, J., Littman, M. L., & Cushman, F. (2017). Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition*. <https://doi.org/10.1016/j.cognition.2017.03.006>
- Huys, Q. J. M., Eshel, N., O’Nions, E., Sheridan, L., Dayan, P., & Roiser, J. P. (2012). Bonsai Trees in Your Head: How the Pavlovian System Sculpts Goal-Directed Choices by Pruning Decision Trees. *PLOS Comput Biol*, 8(3), e1002410. <https://doi.org/10.1371/journal.pcbi.1002410>
- Jordan, J. J., Hoffman, M., Nowak, M. A., & Rand, D. G. (2016). Uncalculating cooperation is used to signal trustworthiness. *Proceedings of the National Academy of Sciences*, 113(31), 8658–8663. <https://doi.org/10.1073/pnas.1601280113>
- Jowett, B. (1925). The Dialogues of Plato. *Journal of Hellenic Studies*, 45(4), 274.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Macmillan.
- Kahneman, D., & Thaler, R. H. (2006). Anomalies: Utility Maximization and Experienced Utility. *The Journal of Economic Perspectives*, 20(1), 221–234. <https://doi.org/10.1257/089533006776526076>

- Kelman, H. C. (1958). Compliance, Identification, and Internalization: Three Processes of Attitude Change. *The Journal of Conflict Resolution*, 2(1), 51–60.
- Koh, H. H. (1997). Why Do Nations Obey International Law? *The Yale Law Journal*, 106(8), 2599–2659. <https://doi.org/10.2307/797228>
- Krupka, E. L., & Weber, R. A. (2013). Identifying Social Norms Using Coordination Games: Why Does Dictator Game Sharing Vary? *Journal of the European Economic Association*, 11(3), 495–524. <https://doi.org/10.1111/jeea.12006>
- Lichtenstein, S., & Slovic, P. (2006). *The Construction of Preference*. Cambridge University Press.
- López-Pérez, R. (2008). Aversion to norm-breaking: A model. *Games and Economic Behavior*, 64(1), 237–267. <https://doi.org/10.1016/j.geb.2007.10.009>
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. (1986). Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Vol. 1. In D. E. Rumelhart, J. L. McClelland, & C. PDP Research Group (Eds.) (pp. 3–44). Cambridge, MA, USA: MIT Press. Retrieved from <http://dl.acm.org/citation.cfm?id=104279.104284>
- Milgram, S. (1963). Behavioral Study of obedience. *The Journal of Abnormal and Social Psychology*, 67(4), 371–378. <https://doi.org/10.1037/h0040525>
- Neo, W. S., Yu, M., Weber, R. A., & Gonzalez, C. (2013). The effects of time delay in reciprocity games. *Journal of Economic Psychology*, 34, 20–35. <https://doi.org/10.1016/j.joep.2012.11.001>
- Nisbett, R. E., & Cohen, D. (1996). *Culture of honor: The psychology of violence in the South* (Vol. xviii). Boulder, CO, US: Westview Press.
- Nozick, R. (1974). *Anarchy, State, and Utopia*. Basic Books.
- Parsons, T., & Shils, E. (2011). *Toward a general theory of action*. S.I.: Nabu Press.
- Pavlov, I. P., & Anrep, G. V. (1927). *Conditioned Reflexes*. Courier Corporation.
- Phillips, J., & Bloom, P. (under review). Do children believe immoral events are magical?
- Phillips, J., & Cushman, F. (in press). Morality constrains the default representation of what is possible. *Proceedings of the National Academy of Sciences*.
- Phillips, J., & Knobe, J. (2009). Moral Judgments and Intuitions About Freedom. *Psychological Inquiry*, 20(1), 30–36. <https://doi.org/10.1080/10478400902744279>
- Phillips, J., Luguri, J. B., & Knobe, J. (2015). Unifying morality's influence on non-moral judgments: The relevance of alternative possibilities. *Cognition*, 145, 30–42. <https://doi.org/10.1016/j.cognition.2015.08.001>
- Quervain, D. J.-F. de, Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., & Fehr, E. (2004). The Neural Basis of Altruistic Punishment. *Science*, 305(5688), 1254–1258. <https://doi.org/10.1126/science.1100735>
- Raihani, N. J., & McAuliffe, K. (2014). Dictator Game Giving: The Importance of Descriptive versus Injunctive Norms. *PLOS ONE*, 9(12), e113826. <https://doi.org/10.1371/journal.pone.0113826>
- Rand, D. G. (2016). Cooperation, Fast and Slow Meta-Analytic Evidence for a Theory of Social Heuristics and Self-Interested Deliberation. *Psychological Science*, 0956797616654455. <https://doi.org/10.1177/0956797616654455>
- Rand, D. G., & Epstein, Z. G. (2014). Risking Your Life without a Second Thought: Intuitive Decision-Making and Extreme Altruism. *PLOS ONE*, 9(10), e109687. <https://doi.org/10.1371/journal.pone.0109687>
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416), 427–430. <https://doi.org/10.1038/nature11467>
- Rand, D. G., Peysakhovich, A., Kraft-Todd, G. T., Newman, G. E., Wurzbacher, O., Nowak, M. A., & Greene, J. D. (2014). Social heuristics shape intuitive cooperation. *Nature Communications*, 5, 3677. <https://doi.org/10.1038/ncomms4677>
- Raz, J. (1999). *Practical Reason and Norms*. OUP Oxford.
- Rilling, J. K., Gutman, D. A., Zeh, T. R., Pagnoni, G., Berns, G. S., & Kilts, C. D. (2002). A Neural Basis for Social Cooperation. *Neuron*, 35(2), 395–405. [https://doi.org/10.1016/S0896-6273\(02\)00755-9](https://doi.org/10.1016/S0896-6273(02)00755-9)
- Ruff, C. C., & Fehr, E. (2014). The neurobiology of rewards and values in social decision making. *Nature Reviews Neuroscience*, 15(8), 549–562. <https://doi.org/10.1038/nrn3776>
- Samuelson, P. A. (1948). Foundations of Economic Analysis. *Science and Society*, 13(1), 93–95.

- Schauer, F. (1991). *Playing by the Rules: A Philosophical Examination of Rule-Based Decision-Making in Law and in Life*. Clarendon Press.
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A Neural Substrate of Prediction and Reward. *Science*, 275(5306), 1593–1599. <https://doi.org/10.1126/science.275.5306.1593>
- Schulz, J. F., Fischbacher, U., Thöni, C., & Utikal, V. (2014). Affect and fairness: Dictator games under cognitive load. *Journal of Economic Psychology*, 41, 77–87. <https://doi.org/10.1016/j.joep.2012.08.007>
- Sen, A. K. (1997). *Choice, Welfare and Measurement*. Harvard University Press.
- Singh, S., Lewis, R. L., & Barto, A. G. (2009). Where Do Rewards Come From?
- Sunstein, C. R. (1996). Social Norms and Social Roles. *Columbia Law Review*, 96(4), 903–968. <https://doi.org/10.2307/1123430>
- Sutter, M., Kocher, M., & Straub, S. (2003). Bargaining under time pressure in an experimental ultimatum game. *Economics Letters*, 81(3), 341–347. [https://doi.org/10.1016/S0165-1765\(03\)00215-5](https://doi.org/10.1016/S0165-1765(03)00215-5)
- Sutton, R. S., & Barto, A. G. (1998). *Introduction to Reinforcement Learning* (1st ed.). Cambridge, MA, USA: MIT Press.
- Tabibnia, G., & Lieberman, M. D. (2007). Fairness and Cooperation Are Rewarding. *Annals of the New York Academy of Sciences*, 1118(1), 90–101. <https://doi.org/10.1196/annals.1412.001>
- Tetlock, P. E., Kristel, O. V., Beth, S., Green, M. C., & Lerner, J. S. (2000). The psychology of the unthinkable: Taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology*, 78(5), 853–870. <https://doi.org/10.1037/0022-3514.78.5.853>
- Thaler, R. H. (1988). Anomalies: The Ultimatum Game. *The Journal of Economic Perspectives*, 2(4), 195–206.
- Thorndike, E. L. (1927). The Law of Effect. *The American Journal of Psychology*, 39(1/4), 212–222. <https://doi.org/10.2307/1415413>
- Turnbull, C. M. (1987). *The Mountain People* (Touchstone edition). New York: Touchstone.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
- Ubel, P. A. (2001). *Pricing Life: Why It's Time for Health Care Rationing*. MIT Press.
- Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. A. (2011). Does the brain calculate value? *Trends in Cognitive Sciences*, 15(11), 546–554. <https://doi.org/10.1016/j.tics.2011.09.008>
- von Neumann, J., & Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Zaki, J., & Mitchell, J. P. (2011). Equitable decision making is associated with neural markers of intrinsic value. *Proceedings of the National Academy of Sciences*, 108(49), 19761–19766. <https://doi.org/10.1073/pnas.1112324108>