Punishment as communication

Fiery Cushman, Arunima Sarin & Mark Ho

Punishment is traditionally modeled as a “constructed incentive”. Specifically, the person who punishes is constructing a set of incentives that will motivate the person being punished to change her behavior. This model spans moral psychology, behavioral economics, artificial intelligence and comparative psychology. Yet, while it captures something important about the logic of punishment, it also misses important contours of how humans punish and how they learn from punishment. We argue that, in addition to understanding punishment as a form of incentive, we must also understand it as a form of communication. According to the constructed incentive model, a person learns from punishment the way she learns from other negative events: by trying to avoid the state or action that elicited it. In contrast, according to the communication model, a person can learn from punishment by attempting to infer the punisher’s communicative intent. Meanwhile, the punisher can choose punishments that are maximally informative given the learner’s likely inferential stance. This coordinated communicative reasoning enables faster and more accurate learning during punishment and provides a better fit to empirical data on how people actually punish and learn from punishment.

Introduction

Good scientific theories capture a lot of facts with just a few principles. What is the simplest model of human punishment that we can get away with?

For several decades a single, standard theory has looked like a terrific bargain (e.g., Hofmann et al, 2018). This theory begins with the observation that humans (like non-human animals) are motivated to gain pleasure and avoid pain. Thus, we can shape each other’s behavior by constructing incentives: If Alice wishes for Bob to stop littering then she fines him \$5 whenever he does it, and if the pain of the fine outweighs the convenience of littering then Bob will change his behavior.

Over the years, many useful additions have been built upon this foundational model of constructed incentives. For instance, while punishment is designed at the ultimate, “adaptive” level to deter others from doing harm (Boyd & Richerson, 1992; Clutton-Brock & Parker, 1995; Fehr & Gächter, 2002; Henrich & Boyd, 2001), people are motivated at the proximate (i.e.,

mechanistic) level by a relatively blind desire for retribution (Carlsmith & Darley 2000; Weiner 1995). Thus, the “construction” is often a matter of adaptive design rather than intentional planning. Also, the accounting principles that make punishment worthwhile can be quite complex—Alice may punish Bob because of direct benefits to herself (Clutton-Brock & Parker 1995), indirect benefits mediated by group welfare (Fehr & Gächter 2002), reputation (Brandt, Hauert, Sigmund, 2003), her institutional role (Andreoni, 2011; Traulsen, Röhl, & Milinski, 2012), or her desire for competitive advantage over Bob (Raihani & Bshary, 2019). Still, all these theories share a common model of how Alice’s punishment will affect Bob’s behavior: by constructing an incentive contingently linked to some aspect of his behavior.

Is this simple model sufficient? It proposes that the sting of criticism feels more or less like the sting of a nettle. In other words, we learn from social punishments the same way we learn from non-social punishments: Both functions as a kind of reinforcement. But this model seems

to miss something very important about the experience of being punished. When stung by a nettle, people simply avoid nettles. When stung by criticism, however, people don't simply avoid the critic—they try to understand what she meant.

Perhaps, then, even the simplest model of punishment we can get away with requires two parts: incentive and communication. In other words, a person learns from punishment not by mere sensitivity to its incentives, but also by attempting to infer the punisher's communicative intent. Reciprocally, the punisher can choose punishments that are maximally informative given the learner's likely inferential stance. It is in this very natural sense that we would never say a nettle "communicates" to a person when it stings them. Although it certainly shapes the person's behavior, it is not structured around expressed and inferred communicative intent. In contrast, ordinary social punishments like criticism usually are. Punishment thus functions not let nettles, but more like language, demonstration, and other forms of human communication.

In fact, the communicative dimension of punishment is so much like language that it can sometimes be tempting to question whether it is anything *more* than language. Suppose, for example, that you forget an important work deadline and your boss gives you a sharp verbal reprimand expressing her frustration and disappointment. Is this an act of punishment with a communicative dimension, or is it merely an act of linguistic communication? We could adopt a restrictive definition of punishment in which "material" consequences are required, and cheap talk doesn't count. On this view, you haven't been punished by your boss until you've been fined, demoted, fired, etc. But this would miss a commonsense and useful notion in which a verbal reprimand functions very much like these other things: It is a response to a transgression, motivated by anger, designed to modify behavior, eliciting negative affect, and implying the threat of consequences even if not imposing any presently. Your boss has certainly engaged

in linguistic communication, but he has also engaged in a (mild) act of punishment.

To see why this broader definition of punishment is so useful—and the concept of "communication" so indispensable—now consider a slightly different case. This time your boss says nothing about the blown deadline, but the next day she conspicuously fails to bring in a cake for your birthday, as she ordinarily does for all employees. Now, there is some literal sense in which the boss has imposed a material cost by her punishment: You are short one piece of cake. But to conceptualize her punishment in this restricted manner misses its point entirely. Although your boss has not used language, she is surely communicating with you, and her communicative intent is crystal clear. No less than words, her actions convey a reprimand and an implied threat: "You let me down, I noticed, and if it happens again the consequences are going to go beyond birthday cake." Thus, the true incentive value of the boss's action is not the literal loss of cake, but the rich information it conveys. It is not conveyed literally, nor in this case is it conveyed linguistically, but rather in the interplay between two parties capable of sophisticated inferences regarding communicative intent.

This model of punishment, involving both incentive and communication, makes important predictions about when, where and how punishment works. To choose just one example, humans can easily make complex inferences about communicative intent that are challenging for non-human animals. This may explain why punishment is notably rare in non-human animals (Hammerstein, 2003; Raihani, Thornton, & Bshary, 2012; Stevens, Cushman, & Hauser, 2005; Stevens & Hauser, 2004). It is restricted almost exclusively to physical threats or aggression designed to challenge an ongoing behavior immediately, such as when an animal charges at something that invading its territory. Similar insights apply the way that children punish and learn from punishment, the design of social artificial intelligence, and the design of institutionalized punishments like the criminal justice system.

We therefore have three goals: To distinguish the contributions of incentives

and communication in human punishment; to then show how they are intertwined; and, finally, to show why it matters. Our approach is to consider in close detail two recent studies that define formal models of punishment as communication and that test these models experimentally.

Communicating “action values”

Traditional models of punishment and reward focus exclusively on their incentive value, ignoring their communicative function. One of the reasons such models are appealing is because they make simple and clear predictions grounded in the basic mechanics of reward learning and expected value maximization. Suppose, for instance, that at one time you and your spouse each did the dishes about as often as the other, without much effort or coordination. Neither of you ever bothered to say “thanks” to the other because it was a responsibility that you both understood to be shared. Recently, however, your spouse has begun to do dishes less often, leaving more for you. In fact, over the last couple of weeks, your spouse hasn’t done the dishes even once. But then, tonight, on a whim, your spouse does all the dishes. You might say “thanks so much for doing the dishes!” with enthusiasm.

Our understanding of this exchange is so effortless that it is easy to overlook a striking fact: A simple model of reward learning and expected value maximization predicts that your expression of thanks will teach your spouse to do the dishes *less*. Here is what your spouse should learn:

1. *If I do the dishes half the time (as I used to), I will pay the cost of lots of work, and never get the reward of thanks.*

2. *If I do the dishes once every two weeks (as I just did), I will pay fewer costs, and I will get a thank-you every two weeks.*

Thus, your “constructed incentive” would actually incentivize the very opposite of the behavior that you intend.

Computer scientists have encountered precisely this dilemma when attempting to

design basic forms of artificial intelligence (AI) that adaptively respond to human rewards and punishments (Isbell, Shelton, Kearns, Singh, & Stone, 2001). In keeping with the constructed incentive model, they assumed that human evaluative feedback was naturally designed to motivate a reward maximizing agent. Humans, they assumed, would punish the computer for doing the wrong things (e.g., sending the wrong email to the SPAM filter, or retrieving the wrong product from the warehouse) and reward it for doing the right things. Thus, if they designed the AI to maximize human rewards and minimize human punishments, the AI systems would eventually learn to do exactly what humans want. In reality, however, AI systems trained on human evaluative feedback often fail dramatically. This is not because the systems are bad at maximizing reward and minimizing punishment. Rather, it is because the human rewards and punishments, if properly maximized, would not incentivize the behavior that the humans intend. Humans are not constructing appropriate incentive schemes with this punishments and rewards; they are doing something else.

But what? If humans are not using reward and punishment to construct the kind of straightforward incentive that would appropriately shape the behavior of a reward-maximizing agent, what *are* they doing? Is it an error, or does it reflect an alternative, coherent principle of design?

In a recent study taking up this question (Ho et al., 2019), participants were presented with a simple task (Figure 1). Their job was to train an agent (graphically represented either as a dog or as a child) to walk along a path towards a goal without going off the path. For instance, participants assigned to the digital dog had to train it to walk along a stone path to the door of a house without stepping on flowers along the way. In order to do this, participants watched as the dog took various actions, and they then used punishments and rewards to try to train the dog to behave better. (The dog’s behavior was controlled by programs varying in design and complexity across experiments).

Figure 1. The tasks used by Ho et al 2019. In (a), the participant is asked to reward and punish the actions of a dog. The goal is to teach the dog to walk along the path and into the door without stepping on flowers. In (b) the task is identical except that the agent is a person, the path is made of tiles and leads to a bathtub, and the area to be avoided is a rug. The target policy (c) is identical for both tasks.

How did participants in this task use punishments and rewards to try to train the dog? Nearly all participants used one of two strategies. One of these, the “action signaling” method (Figure 2a), gave the dog a reward every time it took a good action: One that moved it closer to the door, along the path. And, it gave the dog a punishment every time it took a bad action: Any action that moved it onto flowers when some other option was available. The other strategy, “state training” (Figure 2b), was nearly identical in concept. Instead of rewarding and punishing “actions”, however, it rewarded and punished “states”: Anything that moved into a permissible state was rewarded, and anything that moved into an impermissible state was punished.

These strategies are quite intuitive. At first blush you might suppose that they are good “constructed incentives” for the dog. But they are not: Counterintuitively, reward-maximizing agents learn exactly the wrong behavior from the human “action signaling” and “state signaling” strategies. These reward maximizing agents are built in an extremely simple way, embodying a form of learning that is widely used in computer science (Sutton & Barto 1999), and fundamental to theories of human behavior in psychology and neuroscience (Dolan &

Dayan 2013). Specifically, the agents attempt to perform the precise series of actions that will earn them the greatest amount of reinforcement: Rewards gained, and punishments avoided. When such a reward-maximizing agent is trained according to humans’ “action signaling” or “state training” strategies, what they learn is to run along the path just until they get to the door (thus gaining a lot of praise), and then to run back along the path (or perhaps by a straighter route through the flowers) to the beginning again so that they can re-experience the rewards of moving through the door. Given an infinite amount of time, these agents would happily repeat the loop an infinite number of times. These loops are often called “positive reward cycles” (Ng, Harada, & Russell, 1999).

In other words, like a crazed dish-avoidant spouse, they maximize reward by trying to perpetuate a training regime of praise. They specifically *avoid* attaining the goal of the teacher because this would end the positive reinforcement of training. And, they don’t mind enduring a bit of punishment if it turns out to be the quickest way back to the beginning of the training regime. (In fact, it may be just what is *necessary* to re-initiate training).

Figure 2. A schematic representation of how participants punished and rewarded various actions by the agents (dog or child) in Ho et al (2019). Arrows represent the average amount of punishment and reward; blue arrows represent averages with positive value (rewards) and red arrows represent averages with negative value (punishments). The length of the arrow is proportional to the magnitude of the absolute value. The direction of the arrow indicates the action in question (i.e., movement from one cell to another). A hierarchical clustering analysis identified two clusters of participant responses. One of these, which the authors interpret as “action signaling”, involves rewarding actions that are in the target policy and punishing actions that are not. The other of these, which the authors interpret as “state training”, involves rewarding actions that terminate in “permissible” squares and punishing actions that do not.

This peculiar failure of the reward learning agent is not an intrinsic result of the experiment or task, but rather a result of the specific way that humans naturally mete out reward and punishment. It is easy to define a set of rewards and punishments that would appropriately incentivize a reinforcement maximizing agent (Devlin & Kudenko, 2012; Ng et al., 1999). For instance, if our participants rewarded the dog only when it entered the door and continued to punish it whenever it stepped on the flowers, a reinforcement maximizing agent would quickly learn exactly what to do.

Thus, humans could easily use rewards and punishments as constructed incentives if they wanted to, appropriately guiding the behavior of reinforcement maximizing agents. But they do not. Instead, what the participants in this experiment seemed to be doing was to use punishment and reward as a channel of *communication* to be interpreted by the learner. One group of participants used rewards and punishments to express which actions were “good” or “bad” to perform (“action signaling”), while another group of participants used rewards and punishments to express which states (i.e., positions in the garden) were “good” or “bad” to occupy (“state signaling”). Presumably the same is true of a spouse who says “thanks” to her derelict dish-doer—she means this not as an incentive to be maximized, but as an expression of values to be interpreted and understood.

Why use rewards and punishments in this communicative way, and not just as a constructed incentive? The added communicative dimension may enable the “teacher” to convey information to the “learner”—and thus to shape their behavior—with greater speed and accuracy

Figure 3. Nearly all subjects (36/39) tested by Ho and colleagues (2019) generated a set of rewards and punishments containing at least one positive reward cycle. The set of positive reward cycles generated by participants is diagrammed alongside the number of participants who generated each one.

than would otherwise be possible. In order to see why, it helps to construe the process of teaching and learning in somewhat more formal terms (see Ho et al 2017).

Let us suppose that the ultimate goal of reward and punishment is for a teacher to guide the behavior of a learner. Call this the learner's "policy"—a set of instructions that tells the learner what to do in any given situation. More formally, then, the policy is a mapping from "states" that the learner can occupy (e.g., positions in the garden) to "actions" that she can undertake (e.g., moving in any of the cardinal directions). The teacher has some target policy in mind (e.g., move north and then east along the path, avoiding the flowers, and finally entering the door), and wishes to convey this target policy as quickly as possible to the learner. Further, assume that the learner is ultimately (i.e., adaptively) incentivized to learn this target policy as well. This may be because the teacher will enforce the policy by fiat, or because the policy is actually useful for the learner (for instance, in the case that the teacher is a parent or mentor).

Given this basic setup, should the teacher and learner coordinate on a scheme where the teacher constructs incentives and the learner maximizes this reinforcement, or should the teacher and learner coordinate on a scheme where the teacher uses reward and punishment as a communicative channel that is interpreted by the learner?

Coordinating on constructed incentives necessitates one of two inefficiencies. On the one hand, the teacher could choose a very simple reinforcement scheme like the one we described above: Reward only the correct final action (*entering the door*) and punish all impermissible actions (*stepping on the flowers*). This scheme appropriately incentivizes the learner, but it is inefficient because it may take the learner a great deal of exploration to eventually discover the one, final action that gets rewarded. It is akin to a tennis coach who teaches the novice player how to play by rewarding them only when they win a match. What the learner wants, of course, is more detailed information about the necessary steps to take in order to win.

On the other hand, as we have seen, if the teacher naively begins to reward various

intermediate steps (e.g., rewarding steps towards the door along the path; or rewarding good serves, firm backhands, proper footwork, etc.), this can lead to positive feedback loops that will ultimately lead a pure reinforcement-maximizing agent astray. It is possible to very carefully construct a set of rewards and punishments that incentivize intermediate steps without unintended consequences (Ng et al., 1999). This is sometimes called a "shaping policy". But engineering and implementing such a policy takes a great deal of effort, as it requires carefully balancing the rewards and punishments that a learner could obtain and recalibrating whenever a learner has been led astray by intermediate rewards.

In sum, leveraging another agent's capacity for reward learning feels like a quick and easy way to shape behavior. In fact, however, it sets up an inefficient competition between the teacher and learner—who, rather than adopting the teacher's value function, seeks to exploit the rewards of the lessons themselves. In contrast, coordinating on a communicative scheme like "action signaling" (reward "good actions", punish "bad actions") is highly efficient for both the teacher and the learner. Every time the learner takes any action—whether intermediate or final—the teacher has an opportunity to convey whether it belongs to the target policy or not. And, it is extremely cognitively efficient for the teacher to determine what to reward and what to punish. All she needs to do is compare the learner's actual behavior against the intended target policy.

It therefore makes sense that humans do not appear to actually structure rewards and punishments in the manner appropriate to reward maximization, but instead in a manner that allows for the efficient communication of value functions. This approach to teaching seems to be relatively inflexible: In the studies reported by Ho et al. (2019), participants who trained reward-maximizing agents generally failed to adapt when their feedback was exploited, suggesting they had a strong bias against using their rewards as literal incentives.

Communication by recursive mental state inference

How do people actually implement “action signaling”? In principle, the psychological mechanisms could be extremely simple. Properly coordinated, the teacher simply rewards acts in the target policy and punishes acts absent from it; the learner simply encodes reward acts in new policy and excludes punished ones. In practice, however, such a simple, hard-coded mechanism is both unlikely and suboptimal.

To see why, it helps to begin with a concrete case. Envision another dishwashing dispute, but this time between roommates. Alice and Bob are each supposed to do all their dishes at the end of every meal, but Bob has a growing habit of leaving his dirty dishes around the kitchen for days on end. So, Alice goes to the hardware store, picks up a brand-new sponge and dish soap, and leaves these on Bob’s bed with a note that says: “Love, your roommate”.

This would obviously fail as a mere constructed incentive: Bob is getting rewarded for his dereliction, and a reward-maximizing agent would learn to keep avoiding dishes. But notice that Alice’s strategy would also fail as an action signal according to the simple scheme defined above. Because she has given a gift—a reward—Bob should infer that his recent behavior of leaving dirty dishes around belongs to Alice’s target policy.

Yet, when people are presented with descriptions of this case (and others like it), they naturally predict that Alice’s actions are likely to correct Bob’s behavior, causing him to do the dishes more often in the future (Sarin et al, in prep). Specifically, they endorse the idea that Bob is going to “Get Alice’s message”, understanding the punitive communicative intent behind her ironic gift. Even though in some objective sense Alice gave Bob a reward, they say that Bob will feel bad because of the message Alice intends to convey. Intuitively these claims make sense, but what cognitive mechanism is necessary to explain them?

We posit that humans interpret punishment and reward not via simple action signaling mechanisms alone, but rather through a process of *recursive mental*

state inference. (It is recursive in the sense that each party is drawing mental state inferences about the other parties’ mental state inferences). According to this model, learners are attempting to infer the most likely target policy of the teacher given the action (reward or punishment) that the teacher has performed. Anticipating this fact, teachers choose actions that they think will be appropriately interpreted by learners. This model implies several layers of embedded mental state inference:

Learner: What actions get me the most reward and least punishment?

Teacher: What reward or punishment will shape the learner’s behavior towards my target policy?

Learner: What target policy does the teacher intend to shape?

Teacher: What target policy will the learner infer?

Learner: What inference is the teacher trying to get me to draw about her communicative intent?

Embedding this logic in a specific case, Bob asks: “What is Alice thinking I’m going to infer about my recent behavior, based on this unexpected gift?”. Or, your spouse asks, “What is my partner thinking I’m going to conclude when she says, “Thanks for doing the dishes”?”

We know that humans have a remarkable capacity for recursive mental state inference, and it is fundamental to successful communication in other domains. For instance, when people interpret “communicative demonstrations”, such as a knowledgeable adult demonstrating how to use a new tool to a naive child, they sometimes rely on recursive mental state inference (Ho et al., 2018). The adult may perform certainly exaggerated or highly diagnostic actions for the benefit of the child, knowing that the child will draw easier and more accurate inferences by asking herself, “what is the adult trying to convey?”.

Similarly, recursive mental state inference seems to play an important role in structuring linguistic communication in humans (Frank & Goodman 2012). A vivid example is “figurative speech”, such as irony

(Kao & Goodman 2015). Suppose that there is a thunderstorm and your friend says, “Great weather, huh?”—what should you conclude about his preferences? One possibility is that he is being literal and loves thunderstorms; a more likely possibility is that he is being ironic and hates them. We may arrive at the non-literal interpretation of many statements by recursive mentalizing, asking ourselves, “What did my friend think I would interpret this statement to imply about his preferences?”

Thus, the ironic gift of a sponge and soap can be understood in a manner similar to the ironic statement about the storm. Although the sponge and the soap are rewards at a “literal” level, they can attain a different “figurative” meaning when the learner infers the teacher’s communicative intent. Potentially, then, the same general cognitive mechanisms that support inferences about communicative intent for demonstration and language also support inferences about communicative intent in cases of punishment and reward.

Of course, most punishments are not ironic. Neither are most things we say, and yet ironic speech teaches us important lessons about the organization of language. Similarly, ironic punishment is not an interesting case because it’s common, but because it reveals a basic architectural principle of how we ordinarily learn from punishment. It shows that we naturally respond to punishment by trying to understand the communicative intent of the punisher.

In the previous section we described “action signaling”, a coordinated method of teaching and learning from punishment that is potentially much simpler than recursive mentalizing. According to the simplest version of the action signaling model, all that the teacher does is to reward actions that belong to the target policy and to punish ones that do not, and all that the learner does is to update their policy accordingly. So far, we have seen some evidence that this simple action signaling mechanism isn’t sufficient to account for certain patterns in how people respond to reward and punishment (for instance, in cases of ironic punishment). But is there some advantage of recursive

mentalizing that explains why people employ it, rather than a simpler mechanism?

It makes sense for a learner to attempt to understand the communicative intent of a teacher because “local” bits of information about the target policy can afford useful “global” inferences that allow learning to proceed more quickly and reliably (Ho et al., 2017). Consider again the experimental setup employed by Ho and colleagues (Figure 1) in which a teacher tries to train a simulated agent to walk along the path to the door while avoiding the flowerbed. What happens if the agent occupies the lower left cell (i.e., the first step of the path) and gets punished when it tries to walk on the flowers? If it employs a simple action signaling update rule then it will conclude nothing more than that this single, local action does not belong in its target policy.

If, instead, it attempts to infer the communicative intent of the punisher—and if it brings to bear a wide array of plausible background information—then it would be in a much better position to draw many useful inferences:

1. Given that she does not want me to walk there, perhaps she wants those flowers untrampled.
2. If she wants those flowers untrampled she probably wants her other flowers untrampled.
3. If she wants her flowers untrampled she probably also wants them unpicked.

And so on. Similarly, if the teacher rewards the learner for moving along the first step of the path, the learner might reasonably infer that the teacher’s intended policy is for her to follow the path for some further distance, and perhaps to its salient endpoint: the door.

These particular inferences do not necessitate *recursive* mentalizing but, rather, more limited inferences about communicative intent. Nevertheless, they illustrate how “action signals” can be elaborated into a more inferentially rich and expressive form of communication via inferences about communicative intent. To the extent that learners interpret action signals in this manner, teachers can exploit that by choosing the particular signals that

maximize the likelihood of successful learning; learners can exploit this by reasoning about the teacher's principles of maximum informativeness. These elements comprise a basic model of recursive mentalizing.

The codependence of incentive and communication

The standard "constructed incentive" model of punishment is too simple for practical purposes. Rather, in order to approximate even the most basic contours of how people punish and learn from punishment, it is necessary to model punishment as a communicative act intended to convey a target policy. How do these two elements of punishment—*incentive* and *communication*—operate in tandem?

To begin with, in the absence of an actual or implied incentive, mere communication would presumably often lack any motivating power—it would be nothing more than "cheap talk". This would be problematic specifically when the initial incentives of the teacher and the learner are misaligned. In these cases the teacher is using punishment to compel the learner to act in a way she would not otherwise choose. Such cases are common. Punishing theft, for instance, is designed to get a "learner" to stop stealing things that she would choose to take in the absence of punishment. The punisher will presumably fail to prevent future theft by merely communicating "I wish you would stop doing that..." in the absence of any implied "... or else."

Crucially, then, part of the communicative intent of punishment is to imply "... or else": a threat of future punishment in response to future transgression. For instance, when a roommate leaves a new sponge and dish soap on her roommate's pillow, part of the intended message may be, "If you cannot fix this behavior, you're either going to have to find a new roommate or a new apartment". On the assumption that finding a new roommate is costly and undesirable, then, an important part of the communicative act is to transmit information about a constructed incentive. Put simply, communicative punishment

often conveys a threat, and this threat is a form of constructed incentive.

There are, however, many cases in which the incentives of a teacher and learner are actually aligned—where, once the learner understands the communicative intent of a teacher's evaluative feedback, she will be intrinsically incentivized to follow the teacher's target policy. Parents, coaches, teachers, mentors and others often use evaluative feedback not to impose their preferences on a learner's policy, but to help the learner accomplish the goals she already possesses. In these cases, evaluative feedback need not carry an implied "... or else."

There is a second and quite distinct way in which the incentive value of punishment enables effective communication. It serves, in the language of Schelling (1960), as a kind of "focal point" that coordinates the communicative intent of the teacher with the inference of the learner. In principle, if your roommate is failing to do the dishes you could attempt to communicate this by leaving a blank post-it note on his bicycle. This would be a successful act of communication if your roommate coordinated spontaneously on a common understanding of your communicative intent—in other words, if the two of you shared the common prior assumption that blank post-it notes on bicycles often imply dish-doing dissatisfaction. But this, of course, is unlikely. Instead, it is necessary to choose a punitive act that serves as a coordinating "focal point" because both parties share the prior expectation that it might reasonably be chosen as a punishment for the transgression in question. One way to do this is to choose an act that is semantically related to the transgression; for instance, leaving a new sponge and dish soap on his pillow. Another potential solution is to choose an act that serves as coordinating "focal point" because it imposes a cost on your roommate that cries out for explanation, and the simple "constructed incentive" model of punishment furnishes such an explanation. For instance, you could throw his toothbrush in the toilet. Although not semantically related to the dishes, this act imposes a large cost (i.e., incentive) on your roommate at no apparent benefit to yourself. Thus, it is hard to explain on any

hypothesis other than that you are very upset with your roommate and want to change his behavior. A particular clear message might combine both forms of focal point, semantic and incentive-based: For instance, stacking a pile of dirty dishes on his bedroom pillow.

In sum, teachers and learners are in a better position to coordinate on the idea that a punishment is designed to convey dissatisfaction with the learner's behavior precisely because punishment will often have the effect of discouraging it. This common understanding of punishment's role as a disincentive can usefully initiate the process of inferring the precise target policy that the teacher intends to convey.

Conclusion

The simplest model of punishment posits that a teacher constructs a system of incentives which, once maximized by a learner, causes the learner to act in the way the teacher intends. In other words, it describes punishment as a “constructed incentive”. This model is not wrong—indeed, it captures arguably the most essential feature of punishment. Still, it misses too much to be satisfactory by itself. In addition to incentivizing behavior, punishment is used as a form of communication. This structures punishment in such fundamental ways that often punishment will not usefully shape the behavior of a purely “reinforcement-maximizing” agent at all. Rather, the logic of punishment more closely resembles other forms of teaching and social learning, such as teaching and learning from examples (Shafto, Goodman, & Griffiths, 2014) or demonstrations (Ho et al., 2018). In particular, evaluative feedback reflects a kind of “action signal”—an intentionally informative commentary on what to do and what not to do. Learning in this manner becomes especially efficient when the teacher and learner are coordinated in a scheme of expressing and inferring communicative intent via recursive mentalizing.

Acknowledgement

This work was supported by grant 61061 from the John Templeton Foundation.

References

- Andreoni, J. (2011). Gun For Hire: Does Delegated Enforcement Crowd out Peer Punishment in Giving to Public Goods? 1–22.
- Boyd, R., & Richerson, P. (1992). Punishment allows the evolution of cooperation (or anything else) in sizeable groups. *Ethology and Sociobiology*, 13(3), 171–195.
- Brandt, H., Hauert, C., & Sigmund, K. (2003). Punishment and reputation in spatial public goods games. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1519), 1099–1104.
- Bryan, J. H., & London, P. (1970). Altruistic behavior by children. *Psychological Bulletin*, 73(3), 200–211.
- Carlsmith, K. M., Darley, J. M., & Robinson, P. H. (2002). Why do we punish? Deterrence and just deserts as motives for punishment. *Journal of personality and social psychology*, 83(2), 284.
- Clutton-Brock, T. H., & Parker, G. A. (1995). Punishment in animal societies. *Nature*, 373(6511), 209.
- Devlin, S., & Kudenko, D. (2012). Dynamic potential-based reward shaping. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume1* (pp. 433–440). International Foundation for Autonomous Agents and Multiagent Systems.
- Dolan, R. J. and P. Dayan (2013). Goals and habits in the brain. *Neuron* 80(2), 312–325.
- Fehr, E., & Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415(6868), 137.
- Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998–998.
- Grusec, J. E., & Redler, E. (1980). Attribution, reinforcement, and

- altruism: A developmental analysis. *Developmental Psychology*, 16(5), 525.
- Hammerstein, P. (Ed.). (2003). *Genetic and cultural evolution of cooperation*. MIT press.
- Henrich, J., & Boyd, R. (2001). Why People Punish: Defectors Weak Conformist Transmission Can Stabilize Costly Enforcement of Norms in Cooperative Dilemmas. *Journal of Theoretical Biology*, 208(1), 79–89. doi:10.1006/jtbi.2000.2202
- Ho, Mark; Cushman, F.; Littman, M. & Austerweil, J. (2019) People Teach with Rewards and Punishments as Communication not Reinforcements. *Journal of Experimental Psychology: General* 148(3), 520
- Ho, M., Littman, M., Cushman, F.A. & Austerweil, J. (2018). Effectively learning from pedagogical demonstrations *Proceedings of the Cognitive Science Society*.
- 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*, 167, 91-106.
- Hofmann, W., Brandt, M. J., Wisneski, D. C., Rockenbach, B., & Skitka, L. J. (2018). Moral punishment in everyday life. *Personality and Social Psychology Bulletin*, 44(12), 1697-1711.
- Kao, J. T., & Goodman, N. D. (2015). Let's talk (ironically) about the weather: Modeling verbal irony. In *CogSci*.
- Maccoby, E. E. (1992). The role of parents in the socialization of children: An historical overview. *Developmental Psychology*, 28(6), 1006–1017.
- Ng, A. Y., Harada, D., & Russell, S. (1999). Policy invariance under reward transformations: Theory and application to reward shaping. In *Proc. 16th (Vol. 99, pp. 278–287)*.
- Raihani, N. J., Thornton, A., & Bshary, R. (2012). Punishment and cooperation in nature. *Trends in ecology & evolution*, 27(5), 288-295.
- Rihani, N & Bshary, R., (in prep). Punishment: one tool many uses.
- Sarin, A., Ho, M., Martin, J., Cushman, F. (in prep) Punishment and recursive mentalizing.
- Schelling, Thomas C. (1960). *The strategy of conflict* (First ed.). Cambridge: Harvard University Press.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT press.
- Stevens, J. R., Cushman, F. A., & Hauser, M. D. (2005). Evolving the psychological mechanisms for cooperation. *Annu. Rev. Ecol. Evol. Syst.*, 36, 499-518.
- Stevens, J. R., & Hauser, M. D. (2004). Why be nice? Psychological constraints on the evolution of cooperation. *Trends in cognitive sciences*, 8(2), 60-65.
- Traulsen, A., Röhl, T., & Milinski, M. (2012). An economic experiment reveals that humans prefer pool punishment to maintain the commons. *Proceedings of the Royal Society of London. Series B: Biological Sciences*.
- Weiner, B. (1995). *Judgments of responsibility : A foundation for a theory of social conduct* (pp. xvi–301 p.). New York: Guilford Press.