# The Role of Social Reinforcement in Norm Transmission and Cultural Evolution

## *Introduction*

The main claim of this paper is that social forms of reinforcement play a role in cultural evolution. More specifically, I argue that starting early in human history, social approval and disapproval of behavior functioned as reinforcement and punishment; approval and disapproval thereby start a process whereby norms are transmitted (I'll understand norms in a minimal sense: as a regularity in behavior. See §2). The importance of the main claim is that evolutionary theorists have not said much about the role of reinforcement in cultural evolution.[1] Here, I consider four reasons we have for expecting that reinforcement plays a role in cultural evolution, and then I provide a brief layout of the rest of the paper.

First, since reinforcement learning—i.e., the process by which certain events (called rewards and punishers) influence the probability that a behavior will reoccur—is found in various non-human organisms, we have reason to believe that it is an old system. So, we can expect that reinforcement was present from the start of hominin evolution. And the earlier reinforcement appears in history, the more time evolution has to creatively work and innovate. Second, we have strong theoretical grounds for expecting that reinforcement plays a role in cultural evolution. One source of evidence is the various models of the evolution of signaling systems that assume reinforcement (see LaCroix, 2019; Skyrms, 2010 for an overview). Another source of evidence comes from the success of "reinforcement" as a concept in neuroscience, psychology, engineering, and related fields. We should be surprised if reinforcement didn't play a role in cultural evolution.[2]

---

[1] An exception is work by Heyes (2023). I am primarily referring to work by Boyd, Richerson, and Henrich.

[2] Boyd and Richerson (1985) considered reinforcement only as a form of individual learning and contrasted it with cultural learning. This paper suggests that reinforcement plays a role in cultural learning. I'll discuss this point in §2.

Third, if we can appeal to reinforcement to explain the transmission of cultural norms, then we can see one way by which the cultural evolution of norms begins and how cultural learning increases in importance without relying on special adaptations dedicated to or evolved for norm transmission. Various authors hold that humans have genetically based adaptations specialized for culture (e.g., Richardson & Boyd 2004; Henrich 2016; Tomasello 2014). A main challenge for this approach is to explain how cultural learning becomes important enough for evolution to select for genetically inherited capacities for cultural learning. My claim is that reinforcement can explain the transmission of norms without relying on specialized adaptations.[3]

Fourth, Birch (2021) argues that the evolution of normative cognition began with cognitive adaptations for representing standards in practice and skill. [4] Reinforcement can supplement this account if reinforcement can be used to explain standards implicit in human practice and skill. So, to summarize, we are confident that reinforcement was an early evolutionary resource. Appealing to reinforcement in cultural evolutionary theory has strong theoretical support, and doing so has various theoretical benefits.

I will now outline the rest of the paper. §1 sets the context of the paper. In §2, I defend my main claim, namely, that social forms of reinforcement play a role in cultural evolution. I provide a how-possibly account whereby approval and disapproval function as reinforcers and punishers, and this allows for early normative behavior. I also propose a process by which we become reinforceable by approval and disapproval, and I suggest that this process can be the target of selection. In §3, I explain how my account of social reinforcement can explain an important kind of rule based social learning strategy; I thereby make a close connection between

[3] Of course, this is consistent with the view that as cultural learning becomes increasingly important, evolution selects for such specialized adaptations.
[4] Also, see Railton (2021)

these social learning strategies and cultural evolution. As an example, I introduce the case of divination practices.

## §1. A Brief Discussion on Cultural Evolution

I will work within what is commonly referred to as the "California school." This includes the work of Boyd, Richerson, Henrich and their collaborators. I don't give any argument that this is the *correct or complete* characterization of the theory of "cultural evolution." I am simply interested in working out a view from this perspective and in contributing to this tradition.

How can I be said to contribute to this tradition? I wrote in the introduction that we have strong reasons to expect reinforcement to play a role in cultural evolution, but its role is underdeveloped (or neglected) in the California school. One of the reasons that reinforcement may not play a larger role in their theories is that reinforcement is a typically understood as a mechanism of individual learning; it is thereby seen as a costly form of learning. A key theme of Boyd and Richerson's (1985) formal models is that different forms of learning (e.g., individual vs cultural) have different associated costs. They thereby conclude that evolution selects for different forms of learning. It selects for cultural learning when it is less costly than individual learning. In §2, I argue that social forms of reinforcement can be sufficiently low cost. Following Boyd and Richerson's underlying logic, I conclude that evolution would have selected for these social forms of learning.[5]

[5] I will leave open the question of how my social reinforcement account relates to other forms of learning proposed by members of the California school. For example, how does reinforcement relate to context bias or payoff-biased transmission (Kendal et al., 2009)? How do social forms of reinforcement relate to the models predicting the evolution of cultural learning? Can social forms of reinforcement help avoid commitment to evolutionary psychology (Heyes, 2018a)? In this paper, the main contribution is to show that social forms of reinforcement can play a role in cultural evolution.

Lastly, I should note the differences between my account and that of Castro et al's work on the role of approval in cultural evolution (2010, 2022). I agree with these authors that approval and disapproval constitute a cultural inheritance system, that reinforcement is the framework to understand how approval and disapproval function in cultural evolution, and that such account helps explain the origins of normativity (Castro & Toro, 2022; Castro et al., 2021).[6] But here provide distinct arguments for these claims.

§1.1 General Aspects of Selection: Variation, Replication, and Differential Success

In what follows, I will focus on three aspects of evolution by natural selection commonly known as Lewontin (1970) conditions: variation, transmission, and differential success. It's important to note that the discussion in this section is not meant argue that cultural evolution is selectionist. I am using these concepts to frame the discussion. I will say about more about cultural evolution and selection in §2.2.

§1.2. Different views on Transmission, Variation, and Differential Success

Within cultural evolutionary theory, theorists differ on what causes and constitutes variation, transmission, and differential success. I start with the question, "what varies?" The cultural entities that Henrich (2016) discusses include practices, norms, and representations. Heyes (2018a) argues that cultural evolution can also explain cognitive capacities, not just cultural products.[8] Given a genetic starting kit (social motivation and tolerance, low specificity attentional biases, and domain general cognitive processes, e.g., associative learning), Heyes argues that we can explain the development of various cognitive capacities such as imitation,

[6] I do not commit to their suggestion that humans had to evolve a specialized categorization capacity. While I cannot argue that here, I think reinforcement learning itself can constitute categorization capacities. (Radulescu et al., 2019; Aly et al. 2022).

[8] Heyes does not identify as part of the California school.

mind reading, and language.[9] My account is like Henrich's in that I focus on norms. However, my account is like Heyes's in that it appeals to domain a general mechanism, i.e., reinforcement. Furthermore, my account is unlike Henrich's (2016) in that I don't require a cognitive mechanism dedicated and evolved for cultural learning. Henrich has argued that we are genetically predisposed to be norm following creatures. My view proposes that norm following begins with reinforcement.

Next, what are the causes of variation in cultural evolution? Two sources of cultural variation are imperfect transmission of cultural information and haphazard discovery (Henrich & Boyd 2002). For example, Henrich (2016) writes that given enough time, an individual can accidently discover a slightly new but useful variation on a practice or tool use. By multiplying the *number* of individuals and the *amount* of time, we greatly increase the probability of a member of a population "coming across" advantageous practices and tools. Similarly, Sterelny (2012) appeals to populational and demographic features.[10] A larger population is better able to support specialists. Specialists are more likely to innovate. I draw on these ideas and apply them to norms. Useful norms are discovered accidentally and on the basis of demography.

Besides differences regarding what varies, theorists disagree on how culture is transmitted. One mainstream view, championed by Richardson and Boyd (2004) and Henrich (2016), posits genetically inherited learning tendencies that predispose an individual to copy some agents more than others.[11] These include preferentially attending to cues related to prestige, success, sex, or age. In §2, I argue that reinforcement can explain the transmission of

[9] See Buskell (2021) for an analysis of Heyes and a proposal of a view along similar lines. See Nichols et al. (2019) for another assessment.
[10] As with Heyes, Sterelny does not identify as part of the California school, but I think both authors make noteworthy contributions such that they cannot be ignored.
[11] Commitment to genetically inherited mechanisms for cultural learning may be less central to their theory.

norms. Lastly, let's turn to differential success. The idea behind differential success is that some cultural entities contribute, and some do not contribute, evolutionary benefits.[12] Those practices that do not contribute eventually "die" out.

*§2 A Social Reinforcement Account of Norm Transmission in Cultural Evolution*

In this section, I argue that social forms of reinforcement can play a role in cultural evolution. The role they play is that approval and disapproval are mechanisms of norm transmission. §2.1 elaborates on this claim and gives two arguments in favor. Next, I suggest that to be good arguments, something like the following must be true. (1) Some evolutionarily beneficial norms can emerge as a byproduct of minimally cooperative group life, and (2) incentives are somewhat aligned in those forms of group life. I close by drawing out reasons for both claims. The account so far is consistent with various explanations of how we come to be reinforced by approval. But in §2.2, I propose a three-part process by which domain general learning mechanisms, along with the appropriate interaction with the social-cultural environment, explain the fact that we are reinforced by approval. Moreover, I suggest that this process can be the target of selection insofar as facts about what is a reinforcer for an organism are the target of selection.

*§2.1 Reinforcement and Norm transmission.*

My account of reinforcement as a mechanism of norm transmission in cultural evolutionary theory is inspired by Baum's (2017, 2019) behaviorist account of rule giving and following. However, I am not a behaviorist. And, for it to apply as far as possible into the evolutionary past, my account does not assume complex forms of language.[13] So, my views will

[12] I don't take a stand on whether it is the organism or the group that benefits.
[13] There is at least one strong reason why we should seek an account of norms that does not presuppose complex forms of language. Namely, language itself depends on norms in at least two senses. Formally, language requires certain norms. Materially, complex forms of communication depend on complex forms of cooperation. In turn,

look different and be argued for differently. I'll think of normative claims as emerging at a proto-linguistic level. Here, agents might express various forms of approbation and disapprobation—e.g., grunts or cries. In turn, these expressions influence the behaviors of others, and this begins the emergence and evolution of norms (or proto-norms). Some might think that, at this level, we have neither "claims" nor something "normative." That is fine. On my account, it is approbation and disapprobation that does causal work—it is what reinforces norms. By a norm, I will just mean a standard or a rule describing a behavioral regularity, and these needn't be linguistically represented.[14] That said, I intend this as an account that can scale up along with complex forms of language. One way this can occur is that language allows agents to make explicit what occurs, what is expected to occur, what one's values and commitments are, and so forth. Another way is that linguistically representation of values allows for new values or forms of valuation because linguistic structures support new inferences.

All of this raises a fundamental question. Why should we expect, from an evolutionary point of view, that approval and disapproval relate to reinforcement and norm transmission? Here I provide two answers. One answer is that approval as a reinforcer is an efficient learning mechanism when compared to learning solely by the consequences of one's actions. The non-social consequences of one's actions can be costly or lethal, such as when one consumes a poisonous mushroom or explores unfamiliar territories. Also, it can be difficult to discern what

---

complex forms of cooperation depend on norms of cooperation. The point is not that norms are in some philosophical sense "pre-conditions" for language. Plausibly, both norms and language interact in mutually supporting, complex, and dynamic ways. Rather, the point is just that an account of norms that presupposes language from the start will seem to raise important questions and not obviously answer them.

[14] Philosophers often think of norms as abstract entities. And they think of norms as having a distinct normative force, in a philosophical sense. I prefer an empirical account that begins by looking at what occurs in a society, e.g., the types of verbal and non-verbal behavior that one finds. From here, it is an open question to what extent norms in the philosophical sense are already present in a community or are the result of reflection and elucidation (e.g, Brandom, 1994). It's also a further question how to make sense of norm-internalization and normative force (in a philosophical sense).

such consequences are because cause and effect don't always occur simultaneously and because various other putative causes will be available. Difficulties associated with learning via non-social consequences can be understood as cost problems. Learning by direct experience that a mushroom is poisonous costs one's life. Figuring out which techniques best help produce crops costs years of one's life. These factors make the consequences of one's actions into imperfect reinforcers and punishers. Now, approval and disapproval regarding our actions can help avoid these costs. First, they reduce the lethality cost. A scolded child may experience emotional pain, but they learn to not to eat a poisonous mushroom. Next, as noted, it is difficult to learn from one's mistakes when the feedback is far removed from one's actions. Approval and disapproval are easily paired with one's actions, so they provide immediate feedback. Suppose that my crop fails next year. Was that because of my technique, the seed, soil, weather, or none of these? A knowledgeable person can simply express their dissatisfaction with my technique as I am performing it.[15]

A central idea in Boyd and Richerson's (1985) cultural evolutionary theory is that, sometimes, evolution will favor shortcuts to individual learning because individual learning, however reliable, comes with costs. Their formal models investigate conditions when evolution will favor forms of learning that avoid these costs. My first argument in favor of approval as a reinforcer is in line with their basic logic, but it does suggest a slight correction. Boyd and Richerson (1985) seem to have too quickly dismissed reinforcement as too costly on the basis that, when used as a mechanism for individual learning, it is too costly. My first argument

[15] For example, a farmer's child may be tasked with putting corn kernels in an entire field and putting them at specific distances from each other. But anyone familiar with children knows what will happen next. Kernels will be dropped. Corners will be cut. Their attention or focus will wander, and they thus make blunders. The disapproval of an adult needs to be nearby and motivating enough to keep the child focused on learning the task.

suggests that reinforcement via social sources is in fact low cost, despite being a mechanism typically used for individual learning.

A second answer for why approval should be reinforcing is that receiving social approval has evolutionarily beneficial outcomes. If so, then evolution probably selected for humans tuned to social approval. One reason that approval is important is because it improves and maintains our relationships. As Henrich puts it:

> …Reputation itself is merely a type of cultural information…. Once our ancestors could learn from each other, say about which foods to eat or how to make a tool, we could also learn from each other about whom not to build a long-term relationship with for activities like hunting, sharing, mating, and raiding… (144)

A good reputation depends on the approval others have towards one's characteristics, personality, tendencies, and so forth. Early human life gave a lot of importance to social capital (e.g., reputation) over material capital since there is less of the latter. In mobile foraging societies, people have fewer possessions because they can carry less as they move. One's reputation might be one's most valuable asset simply because it is one's only asset. Even if a society is no longer mobile, or if it is only partly mobile, their economy hasn't scaled up to the level of modern times; thus, many individuals cannot easily offset losses in social capital by drawing on material wealth. In modern society, money can replace a lot of the need for approval or reputation. A person can pay for childcare, transportation, or food without anyone knowing anything about his or her moral character. But even here there are limits: we easily take our businesses elsewhere if we don't like the way we are treated.

At this point, we should address a possible suspicion that my arguments assume something crucial, such as norms or a mechanism ensuring the reliability of other agents. For example, in

my second argument, we should not imagine an agent entering a vibrant community and needing to be accepted. A vibrant community may have practices of hunting, foraging, resources distribution, and so forth. This certainly is a situation where the approval of others pays off, but that is partly because there are already various norms in place. Where did these norms come from? In a similar manner, we need to be careful about how we imagine my first argument. There, I explained that signals of approval and disapproval meet low-cost requirements of learning. Parents tell children not to eat poisonous mushrooms. Someone scolds me for poor farming technique. A question emerges: how do I know that the farmer wants to teach me instead of deceive me? If our interests are aligned, then I can trust the farmer. What mechanisms ensure that our interests are aligned? At the start, I maintained that one of the attractions of reinforcement-based explanation in cultural evolutionary theory is that reinforcement has been around for a long time. We can thus reasonably expect evolution to have used reinforcement. But if interests must be aligned, then we see some limits of the account. The worry is that we must first have the mechanisms that align the interests of various agents in place before reinforcement begins to work.

It seems that my arguments require something like the following: in very simple group life, (1) minimal, beneficial norms can emerge as the byproduct of agent activity in said form of life; and (2) agents' interests are aligned, at least to some extent. Both of these help alleviate the suspicion that my account assumes that norms already exist or that some mechanism ensuring the reliability of others is already in place. For example, if (1) is true, we can posit some benefits of responding to social approval without assuming that these benefits only occur in the context of a vibrant and fully developed community. If (2) is true, then approval and disapproval can be at least somewhat reliable signals. Aligned interests function as disincentives to deceive someone.

If the farmer in the above example is my father, then he is not incentivized to teach me a bad farming practice. Both (1) and (2) are plausible claims, as we can see by drawing on some work from Sterelny.

Sterelny (2021) argues that hominin lifeways were cooperative as early as 1.8 million years ago.[16] One valuable insight in Sterelny's work involves the distinctions between forms of cooperation and the mechanisms that support these forms. That is, there are different forms of cooperation, different mechanisms supporting these forms, and different ways these interact. Increasing levels of cooperative complexity require increasingly complex mechanisms stabilizing cooperation. Cooperation is destabilized by free riders and those who steal the profits of cooperation, which Sterelny refers to as bullies. The more complex the form of cooperation, the harder it is to deter free riders and bullies. But, Sterelny maintains, early forms of hominin cooperation were simple enough to solve the free rider and bully problem.

One example of said form of early cooperation is collective hunting. In collective hunting, everyone who participates is present, so reciprocation is not required. Anyone who takes more than a fair share will be seen by everyone. Everyone is interested in sanctioning the cheater. So, recruiting third party support to sanction cheaters is not needed. Moreover, Sterelny's account maintains that these early lifeways were egalitarian. Dominance hierarchies, or bullies and free riders, faced the development of weapons. In non-human animal societies, dominance depends on the strength of the individual. Weapons challenge power-based hierarchies. Specifically, weapons play this role when collectives are not yet large and complex enough to form a class with weapons dominating another class.

[16] This estimate is based on evidence of collective hunting. Of course, it is always possible that collective hunting is older than that and that the evidence is either erased or undiscovered.

The relevance of Sterelny's account is that it gives us an early, simple form of cooperation where interests are minimally aligned, and norms begin emerging as a byproduct of simple group life. Consider a simple norm such as, "don't take more than your fair share."[17] If an agent lives in the world as Sterelny describes it, taking more than one's fair share can cause resentment. At best, one is excluded from collective hunting. At worst, one is attacked with weapons. It is certainly in the interest of parents to approve of fair sharing behavior in their children and to disapprove of unfair sharing behavior, and it is in the interest of children to respond to their parent's attitudes. If children learn to share fairly, members of the hunt benefit for at least two reasons. First, the resources were zero-sum. Second, sanctioning group members comes with costs. For example, excluding a member from a collective hunt means there is one less team member. In this way, community members come to have an interest in the early education of children. If they don't intervene directly in children's learning via approval and disapproval, they may intervene indirectly by disapproving of parents who don't themselves intervene in the child's learning. Thus, if Sterelny's account is correct, then we can meet the two requirements. Norm governed lifeways can begin to emerge early in hominin evolution. That is, it can occur as early as the just described basic forms of cooperation evolve. On Sterelny's view, this is almost two million years ago. It is difficult to make judgements about what could have happened in millions of years. But we can put this in perspective by noting that the account I've been developing is plausible at the early stages of hominin evolution. The model also suggests a way forward in cultural evolution. I mentioned above that, as cooperation increases in complexity, we need new mechanisms to solve for cheating and defection. If normative life can

[17] Fairness involves normative concepts, but my account does not presuppose them. As Tomasello (2019) has noted, primates can resent the fact that they get less than they expected. And one way to manage their expectations is to give their peers tasty food and them less tasty food. In my account, "fair share" can mean the somewhat equally distributed reward an agent expects based on some collaborative activity.

begin in minimal forms of cooperation, we get the chance that norms for more complex forms of cooperation evolve.

§2.2 Selection and Reinforcement via Approval

This sub-section develops the idea of a reinforcement profile, and it draws its connection to selection and social-cultural learning. Reinforcement profiles are facts about what is a reinforcer or punisher for an organism. Reinforcement profiles are targets of selection insofar as they vary, are transmitted, and are differentially successful. One way of understanding reinforcement profiles is that they are based on evaluative criteria. Evaluative criteria are standards an organism "uses" for adopting behaviors. Where do these standards come from? A common view is that humans are innately or genetically predisposed to value approval from social sources. While my account in §2.1 is consistent with this view, here I propose an alternative whereby approval is intrinsically valued (as opposed to instrumentally valued), but this value is not innate or genetically inherited. With the help of social sources, the value of approval is learned via secondary reinforcement. According to this proposal, there are rich correlations between the internal states of others (especially caregivers) and effects on oneself. Human, domain general mechanisms learn these correlations. Humans thereby begin using the internal states of others as their own evaluative criteria. The result is a reinforcement profile; humans become reinforceable or punishable by the approval or disapproval of others. Putting this together, insofar as reinforcement profiles are targets of selection, and social reinforcement is the result of a social-cultural process, I suggest that this social form of reinforcement is the result of a cultural-selectionist process.

My first task is to develop the concept of a reinforcement profile. To do so, we will need only a few minimal or core concepts. These minimal concepts are borrowed from the behaviorist

tradition, but nothing here depends on behaviorism (Baum, 2017). The starting point is the assumption that behaviors, in a context, have regular consequences. Some of those consequences make a behavior more likely to reoccur. These are called "reinforcers." Other consequences make the behavior less likely to reoccur. These are called "punishers." To illustrate, each time Maria has asked her grandmother for candy, she has gotten candy. So, the next time she sees her grandmother, she will ask her for candy. Here, her grandmother's proximity is the context. The behavior is asking for candy, and receiving candy is the reinforcer. Lastly, and this is the part that ties back to evolution, the reinforcers are reinforcers because they in general lead to evolutionarily beneficial outcomes, and the punishers are punishers because they in general lead to evolutionarily bad outcomes.

What do I mean by the statement, "reinforcers are reinforcers because they in general lead to evolutionary beneficial outcomes?" To explain what I mean, let me ask a question.[18] Why is candy a reinforcer for Maria? A reinforcer can be broken down into two kinds: proximate and ultimate. The proximate reinforcer is the immediate consequence that reinforces the behavior. It is the sweetness of Maria's candy. The ultimate reinforcer is something that explains why the proximate reinforcer exists. Sugar is a reinforcer for humans because it gives humans energy, hence contributing to their survival (within a limit). Ultimate reinforcers can be identified with Baum's (2017) "HRRR" acronym: health, resources, relationships, and reproduction. So, we can form the question in a more general manner. "Why does candy reinforce Maria's behavior?" can be phrased as "why is candy a proximate reinforcer?" Abstracting even more, "why is a particular consequence a proximate reinforcer for a given organism?" Let's call the facts about

[18] For simplicity, I focus on reinforcers. What I say, with some relevant changes, applies to punishers.

what proximately reinforces an organism's behavior its reinforcement profile. The question could thus be, "why does an organism have the reinforcement profile that it does?"

An evolutionary answer is that different reinforcement profiles will have different consequences—different regarding their HRRR properties. For example, an organism with a reinforcement profile that reinforces food producing behaviors will end up with more energy resources than an organism without it. What this suggests is that reinforcement profiles are the targets of selection. Recall, selection requires three ingredients: variation, replication, differential success. Here, the reinforcement profiles vary, are replicated, and succeed depending on the HRRR-properties of their consequences.

In what follows, I'll understand reinforcement profiles as based on, or describable as, evaluative criteria. Evaluative criteria are rules or standards that function as criteria for adopting or rejecting a behavior. For example, a rat may discover that licking the nozzle of a water bottle produces water. The feeling of satisfaction following the drinking of water can be described as the guiding criteria for future decisions to drink from the water bottle.

How does an organism acquire its reinforcement profile, i.e., its evaluative criteria? A long-standing answer comes from biological evolution: an organism genetically inherits a reinforcement profile that is the result of evolutionary history. Is there an analogue in cultural evolution? I answer affirmatively. Castro and Toro (1995, 2004, 2010) have argued that human social learning contains the ability to acquire evaluative criteria via social sources of approval and disapproval. I am in agreement, but my precise question is, "why do approval and disapproval function as reinforcers and punishers?" If they function as reinforcers and punishers, then approval and disapproval are part of our reinforcement profile. What is the explanation for this fact? Castro and Toro seem to base this reinforcement profile on a specialized adaptation for

social approval. By contrast, I propose a three-part account of how this social aspect of our reinforcement profiles can be acquired without assuming a genetically specialized adaptation.

Here is my three-part outline. First, there are rich correlations between the emotional states of caregivers (or of signs of their emotional states) and the experiences of dependents.[19] Prenatally, a mother's emotional state will correlate with various biochemical states (e.g., dopamine, stress hormones, nutrients, or blood pressure) that will affect the fetus. Postnatally, caregivers are sources of food, touch, kind gestures, warmth, nutrition, and so forth. But when upset, they can be sources of noxious stimuli such as loud verbal sounds. Secondly, early in development, humans have sufficiently robust learning mechanisms–sufficient enough to begin learning about these correlations. Of course, these learning mechanisms take a long time to mature (Sydnor et al., 2021), but the point is that learning begins early even if full maturation comes later. Third, just as humans use their experiences to evaluate events, humans come to use the experiences of their caregivers to evaluate events (Joiner et al., 2017; Borsa et al., 2019). To see this point, consider the following question. How do we learn that chocolate is good and burns are bad? Well, one day we tasted chocolate and that produced a pleasurable sensation. Another day we touched a flame, and it produced a terrible sensation. We came to associate one stimulus with another. As with the chocolate and the fire, the child learns that the emotional states of a parent are correlated with (causally implicated in) their own emotional states. Chocolate and happy parents are good. Fire and angry parents are bad. Why? The former

[19] In my reading, we should think of approval and disapproval as higher order concepts that have their semantic basis on positive and negative experiences. For example, "approval" is an abstract concept. Our understanding of it is likely anchored in early experiences. Approval is correlated with other reinforcers such as touch, gestures, warmth, and nutrition. Disapproval is correlated with, for instance, noxious verbal replies. So, our ability to grasp the concept begins with the kind of biological beings that we are (e.g., we like warmth and dislike sounds of certain pitch ranges), and development and cultural practice help construct the concept.

produces pleasure and the latter displeasure.[20] My proposal is that the value of approval from others can be learned and valued intrinsically.

It will be useful to have a name for the type of reinforcement profile that we have been discussing. Borrowing from Castro and Toro, let's call this type of reinforcement "suadens reinforcement." Suadens comes from the Latin word suadeo, which means to approve or value.

Let's put together the main ideas of this sub-section. The first key idea is that selection can target reinforcement profiles. Different profiles have different HRRR consequences, and selection plays a role here because of these differences. The second key idea of this sub-section is that suadens reinforcement is part of human reinforcement profiles; we acquire suadens reinforcement via social-cultural sources. These two ideas suggest that suadens reinforcement is the result of a (cultural) selectionist process.

To be clear, this last suggestion is contended by various authors who question whether cultural evolution is a genuinely selectionist process (e.g., Chellappoo, 2022; Lewens, 2002, 2015; Sperber, 2000). I mean my proposal—that suadens reinforcement is the result of cultural selection—as a way of understanding the emergence of suadens reinforcement; this is not meant as an argument that would answer all relevant doubts on cultural selection. Those who are skeptical of cultural selection may read the rest of this paper and focus only on the main claims. For example, in the next section, I suggest that social learning strategies are the result of cultural selection insofar as they are accounted by a reinforcement profile that is itself the result of cultural selection. The main claim here is that these learning strategies can be transmitted via

[20] It is important to not over intellectualize this process. We shouldn't think about it as solely occurring at a sophisticated, conscious, and cognitive level. There is now ample evidence suggesting multi-directional pathways between the brain, gut, and microbiota. Much is still unknown about these pathways, but they have been implicated in reward processes, among various other things (Strandwitz, 2018; González-Arancibia et al., 2019; De Wouters et al., 2018; Garcia-Cabrerizo et al., 2021). This suggests that cultural evolutionists may search for very early forms of (non-conscious) social learning, as these processes may occur early in development.

social forms of reinforcement. Readers who are skeptical that cultural evolution is selectionist in this regard can read me as saying "if cultural selection explains our social reinforcement profiles, then it explains social learning strategies. It does so indirectly." Critics can reject the antecedent and still endorse my main claim.

*§3 Evolution, Suadens Reinforcement, and Social Learning Strategies*

In the previous section, I provided a cultural evolutionary account of norms. Here, I propose that this account can help explain rule based social learning strategies. I also propose that these rule based social learning strategies have a cultural selectionist basis: they are the result of cultural selection insofar as they are supported by a reinforcement profile that itself is the result of cultural learning and selection. The value of this contribution can be seen by situating it within the work of Cecilia Heyes, so I will start there. Then, §3.1 provides an example.

There are two broad (non-exhaustive) ways of characterizing Cecilia Heyes' work on cultural evolution. One is the proposal and theory of "cultural evolutionary psychology" or "cognitive gadgets" (Heyes, 2018a, 2023). Another is the idea of culture-culture coevolution (Heyes & Birch, 2021). The former is a research program and theory according to which various human-cognitive mechanisms are the result of cultural evolution. The latter is a model attempting to describe the emergence, development, and distinctiveness of human culture. For both projects, the concepts of *social learning strategies* and *explicit learning biases* are central. I'll briefly outline these concepts and explain how my account can illuminate the evolutionary basis of selective social learning strategies.

Heyes (2016, 2018a) points out that, as researchers use the term, "social learning strategies" refers to a diverse array of learning phenomena that is found in insects, fish, birds,

and primates. The core idea is that an organism's behavior is influenced by the information it gathers from conspecifics. Heyes believes that selectivity makes human social learning strategies distinctive. Selectivity refers to the tendency of acting on or learning from one source of information over the other. There are at least two possible explanations of selectivity. On the strategic view, conscious, explicit, and domain specific rules make human social learning selective. Heyes gives the example of the rule "follow digital natives" (people born during the digital age) when learning about new technology. On another view, called the attentional view, domain general cognitive processes make social learning selective. The idea is that domain general, individual learning (such as that based on associative mechanisms) influences attention. In turn, attention influences what is learned and who one learns from simply because one is paying more attention to some agents over others. Heyes maintains that for the most part, the attentional view is correct. Social learning is selective in both humans and non-humans via domain general, attention directing mechanisms. However, Heyes also thinks that the strategic view is sometimes right; the strategic view explains the distinctiveness of human social learning. These rules make human social learning distinctive because they accumulate the experience of many agents, improve fidelity in learning, and contribute to variation. I'll refer to strategic social learning strategies as SSSs.[33]

It is worth comparing SSSs with explicit learning biases as discussed by Birch and Heyes (2021). We can read these authors as proposing a broader empirical model. Their account begins with the divergence from great apes about 6 million years ago. Early human social learning was based on domain general mechanisms. Agents could thereby acquire knowledge

[33] On my reading, the connections between the attentional and strategic view need to be further explored. For example, what grabs our attention is influenced by culture. If so, the strategic and attentional view may converge in important ways.

and skills from conspecifics (mostly parents). In a later stage, humans began using their cognitive-representational abilities to pass on "explicit learning biases," which are types of human SSSs in our terminology.[34] These have the previously mentioned benefits (e.g., high fidelity transmission), but Birch and Heyes add something, namely, there is "…more to gain from investing effort in copying the specific technique of a specific individual, rather than hedging [one's] bets and learning from as many different models as possible."[35] Explicit learning biases help begin a feedback loop by getting agents to choose between groups. As more successful groups attract more agents, they increase their chances of success, which in turn attracts more agents.

One question emerges from the discussion on SSSs or explicit learning biases, namely, "how are SSSs the result of cultural evolution?" For example, Heyes (2018a, 2018b) does not tell us *how* selective SSSs, which are conscious, explicit, and reportable rules, are the product of cultural evolution. Heyes may be read as inferring from these features *that*, plausibly, they are the result of cultural evolution instead of, say, genetic evolution or design. Even so, it is important to search for a how account (even if it's a how-possibly account) because, otherwise, it isn't entirely obvious how cultural evolutionary *theory* can explain these types of SSSs (Cao, 2020). To illustrate my point, consider Heyes' example of "copy digital natives." It is not obvious why cultural evolutionary theory should be invoked to explain this type of strategy.[36]

---

[34] There are some differences. For example, Heyes seems to think of SSSs as linguistically encoded, or that is the impression given by the examples. By contrast, explicit learning biases needn't be linguistically encoded. Besides this, both play the same role, which is to tell agents who to learn from and how to learn.

[35] This is important because it touches on an often-neglected theme in recent cultural evolutionary literature, namely, that Boyd and Richerson (1985), as well as Henrich and McElreath (2003), place information quality and acquisition costs at the center of their theory.

[36] Laudan (1977) writes that until a theory can explain a phenomenon, it is often unobvious which theory should be able to explain the phenomenon. One of the benefits of my account is that it provides a way for cultural evolutionary theory to explain SSSs. However, I don't think cultural evolutionary theory explains the rule "copy digital natives."

Why do we need cultural evolutionary theory to explain how one learns to use modern technology?

Heyes and Birch's (2021) view is that explicit learning biases play a role in cultural evolution; that is, they form part of the feedback loop mentioned earlier. On their view, one form of cultural fitness (what they call CS2) is understood as the number of models an agent or group acquires. Certainly, this gives SSSs a role to play in cultural evolution, but it leaves open the question of how (if at all) they are the result of evolution or selection. The views I developed in §2.1 and §2.2 make two proposals.[37] First, SSSs are a type of norm that is transmitted via approval and disapproval; they are based on a kind of suadens reinforcement.[38] Second, these learning strategies are the result of evolution or selection insofar as our reinforcement profiles are the result of evolution or selection. In the next sub-section, I will provide a discussion on divination practices as an example of these ideas.

### *§3.1 The Case of Divination*

Heyes (2023) has recently argued for a reevaluation of the idea that humans genetically inherit cognitive and motivational mechanisms specialized for normative life. There are many ways to reevaluate that idea, but the way I do so here is by giving internal norm representation and motivation a *relatively* diminished role. The picture here is one where norms are motivated by the approval and disapproval of social sources. In important ways, the norms are *externally* represented by the reoccurring behavior of others, the practices of a culture, and its traditions.[43] I

[37] The proposal agrees with Heyes (2023) on fundamental details (i.e., reinforcement and norms). The projects, though, focus on different aspects. For example, I am focusing much more narrowly on approval and disapproval.
[38] While I think suadens reinforcement is partially distinct in humans (e.g., other animals can learn to value our approval), I make no claim that it alone makes human social learning distinct. Human learning is distinct due to, among other things, an intricate relationship between memory, mind reading, language, imitation, and domain general learning processes. The work by Heyes and Birch addresses this topic.
[43] This raises philosophical questions about where and how norms are represented. Cultural evolutionary theorists will need to work on this question.

will discuss divination practices as an extended example. I'll say why I understand divination practices as a kind of SSSs, and I argue that these can be explained by social reinforcement.

The earliest hominins inhabited a dynamic and complex world that they knew very little about. They began with only general-purpose learning mechanisms similar to those of other species (see §3). Many animals do quite well with this much. They can learn regularities and predictive relationships between various stimuli. However, these mechanisms are imperfect or insufficient to explain the diversity of human behavior. This is one key insight from Boyd and Richerson (1985) as well as Henrich (2016).

Henrich (2016) conjectures that divination practices evolved to counterbalance our (sometimes imperfect) modes of reasoning.[44] In one example, the Kantus of Kalimantan use bird augury to choose where to sow. Henrich explains that this practice randomizes where one sows, and this randomization is an improvement over relying on one's learning. First, agents have been observed to plant in areas that have recently been flooded because they *mistakenly* believe that a second flood is unlikely. Second, agents are likely to plant where others have had success, and this can lead to bad yields since land nutrients need time to replenish.[45]

Hong and Henrich (2021) understand divination practices as epistemic technologies: means to discover information about the world. Then they ask, "if divination isn't 'real,' then why do divination practices persist across cultures?" In answering this question, they distinguish

[44] Due to limitations of space, I will take Henrich's views as an assumption. But I should note that however strange this idea may appear, it is quite in line with the cultural evolutionary approach of explaining success in terms of evolved cultural practice. Moreover, the point is in line with what various authors, on divergent topics, have claimed (e.g., Sterelny, 2012; Popper, 1979, Novaes, 2012).

[45] The claim is that divination practices can help avoid problems that would result from human reasoning. One interpretation is that reasoning needs help because of cognitive biases, e.g., of the kind that are discussed in psychology. Such biases involve human reasoning operating counter to the prediction of some formal model and doing so under some experimental set up. (I must note that the interpretation of such results as "irrational" is a hotly contested topic). A second interpretation is that even if an agent reasoned according to some good, formal model, they may still get things wrong, and divination can help here. The sowing example is like this in that the prediction about where it will flood is formally correct; the problematic conclusion is merely the result of a small sample.

between objective efficacy (whether divination actually works) from subjective efficacy (whether it's perceived to work). Objective efficacy cannot easily undermine belief in divination because disconfirming evidence is ambiguous. For example, a failed divination might suggest that the ritual was done incorrectly. Several factors explain subjective efficacy. First, biases such as the availability heuristic or saliency effects make it easier to remember times when divination worked. Secondly, the testimony of others, the behavior of practitioners, and the lack of alternative options make even non-believers try divination. In turn, their participation is perceived by others as endorsement.

Altogether, we get the following view: (A) divination practices can have beneficial consequences even if their metaphysical commitments are mistaken. They function to improve action when reasoning may lead one astray. (B) Divination is used as an epistemic tool. (C) Divination practice is sustained by various factors influencing its perceived efficacy. I agree with Hong and Henrich on these points, but I propose an additional factor contributing to "perceived efficacy," and this factor involves a distinct psychological claim. Hong and Henrich interpret the psychology of agents practicing divination as undergoing a cost benefit analysis.[46] Agents engage in divination because they believe it works, or because there are low costs and potential gains if it does work. My psychological explanation, based on the account from §2, is that the practice of divination is grounded on reinforcement principles. This can be a correct explanation even when the same practice is naturally described in cost benefit terms.[47] For

---

[46] Their model is put in Bayesian terminology, but the authors don't commit to its psychological reality. So, the authors seem open to various psychological mechanisms that are consistent with the model.

[47] I think both explanations can co-exist. For example, both cost benefit analysis and reinforcement principles can, at different times, be part of the psychological explanation of engaging divination. Another alternative is that reinforcement is a more fundamental explanation, and cost benefit concepts are ways of making sense of and conceptualizing our decisions—they are ways of making our practices explicit (Brandom, 1994). My argument is open to these or other alternatives.

example, Niv et al. (2002) show that reinforcement learning in bees can lead to risk averse and probability matching behavior.

In broad outline, I propose that reinforcement can play a role at the exploitation-exploration phase. At the exploitation phase, social approval increases the chances that an agent uses divination. At the exploration phase, social disapproval decreases the chances that an agent adopts alternatives to divination. Let me elaborate.

In the exploitation phase, agents engage in divination practices if those practices are sufficiently reinforcing. My claim is that since social forms of approval reinforce divination, approval can go part of the way in explaining divination; that is, they explain why agents engage in the practice (exploitation). Community members will tend to approve of those who consult divine sources. Those who cite divine sources, especially on important matters, will seem reasonable and intelligible. Appearing reasonable and intelligible is a form of receiving social approval since it is a form of social endorsement—one abides by the community's endorsed practices.

How can we be sure that social reinforcement is strong enough to sustain cross cultural divination practices? This question is best answered by formal modeling and empirical investigation. But what has already been said can give us some confidence. Recall Hong and Henrich's insight that disconfirming evidence is ambiguous.[48] This can be rephrased as the claim that, in the exploitation phase, evidence against divination plays a small role in decreasing the chances that one uses divination. If social forms of reinforcement play a significant role for an agent, then it may be strong enough to encourage exploitation.[49]

---

[48] Philosophers of science know this as the problem of underdetermination.

[49] Steven Hassan (2013), an author on cults and a cult-survivor, recommends that one help family members by getting them outside of the cult. Given enough separation, members of cults can come to realize how irrational their beliefs were. The view I am proposing can explain how this works. Cult members get significant validation for

Next, exploration means that the agent searches a given problem space for new behaviors. If exploration is not sufficiently reinforced, then agents won't seek new behaviors. If members of the community discourage variations from a divination practice, then this disincentivizes exploration, i.e., the search for alternatives to divination. If agents do not readily find strongly reinforcing alternatives to divination (e.g., if they don't find something that works better), then social disapprobation of disregarding divination can sustain divination.

I interpret divination practices as a kind of SSSs. For one thing, they are rules about who to learn from, e.g., shamans. Moreover, if Henrich is correct, then divination practices are beneficial to the community even if they are mistaken about its metaphysical commitments. Community members will be interested in *how* one learns from divination practitioners. Some of this practice is directly recommended, as when agents explicitly recommend the practice or shun people for going against it. But it is also indirectly recommended when one sees people observing the practice, following costly or counter intuitive rituals, or collecting various symbolic materials. On this picture, internal representation of norms and internal motivations for following those norms plays a *relatively* diminished role. If my account is correct, then one's culture plays an important role in representing and motivating the norms.

To conclude, the upshot here is that we have an empirically informed case where we can apply my social reinforcement account of norms to understand SSSs. In providing this example of how my account works, I've had to raise (without settling) difficult questions about the

---

holding certain beliefs and engaging in certain practices. Remove the social validation and one is only left with the objective data. In line with this, Williams (2019) argues that belief formation is sensitive to social reinforcers. Williams suggests that this is a form of irrationality, so Williams investigates social-rewards belief formation in cases where it explains irrationality. But I think the approach I am taking here—inspired by Hong and Henrich's epistemic tool model of divination—doesn't say that socially based belief formation is always a form of irrationality, and I don't think Williams' main thesis is logically committed to the irrationality claim. My hypothesis is that social reinforcement mechanisms that sometimes lead to good epistemic practices (e.g., hunting by tracking) can also lead to bad epistemic practices (divination).

evolution and nature of norm representation and their motivation. On my view, these are relatively new and open questions for cultural evolutionary theorists, especially in the California school, and related areas of research (Heyes, 2023).

*Conclusion*

The main goal of this essay is to develop a social reinforcement-based account of norm transmission in cultural evolution. My claim is that humans are sensitive to the approval and disapproval of others, and this plays an important role in the emergence and evolution of norms. To make the account vivid, I argued that it can help explain the evolutionary basis of strategic social learning strategies. I illustrated this with the account of divination. One benefit of this project is that it draws a closer connection between a theoretically successful concept (reinforcement) and cultural evolution. Moreover, the account promises to help us understand the cultural evolution of something important in everyone's life: norms.

# References

Aly M. H. Abdou K. Okubo-Suzuki R. Nomoto M. & Inokuchi K. (2022). Selective engram coreactivation in idling brain inspires implicit learning. *Proceedings of the National Academy of Sciences*. https://doi.org/10.1073/pnas.2201578119

Baum, W. (2017) *Understanding behaviorism: Behavior, culture, and evolution.* Oxford: John Wiley & Sons.

Baum, W. (1995). Rules, Culture, and Fitness. *The Behavior Analyst*, 18, 1-21

Birch, J. (2021). Toolmaking and the evolution of normative cognition. *Biol Philos*, 36, 4. https://doi.org/10.1007/s10539-020-09777-9

Borsa, D. Heess, N. Piot, B. Liu, Siqi. Hasenclever, L. & Munos, R. (2019). Observational Learning by Reinforcement Learning. AAMAS.

Boyd, R. & Richerson, P(1985). *Culture and the evolutionary process*. University of Chicago Press.

Brandom R. (1994). *Making it explicit: reasoning representing and discursive commitment*. Harvard University Press.

Buskell, A. (2021). Cognitive novelties, informational form, and structural-causal explanations. *Synthese*, 198, 8533–8553. https://doi.org/10.1007/s11229-020-02585-4

Cao, R. (2020). Crowding out memetic explanation. *Philosophy of Science*, 87(5), 1160-1171. https://doi.org/10.1086/710518

Castro, L., Castro-Nogueira, L., Castro-Nogueira, M, A. *et al.* (2010). Cultural transmission and social control of human behavior. *Biol Philos* **25**, 347–360. https://doi.org/10.1007/s10539-010-9201

Castro, L., Castro-Nogueira, M A., Villarroel M. & Toro M, A. (2021). Assessor teaching and

the evolution of human morality. *Biological Theory*, 5–15. https://doi.org/10.1007/s13752-020-00362-7

Castro, L & Toro, A, M. (2022). The Origins of Normativity: Assessor Teaching and the Emergence of Norms. *Metode Science Studies Journal.* 13. https://doi.org/10.7203/metode.13.21755

Chellappoo A. (2022). When can cultural selection explain adaptation? *Biology & Philosophy*. https://doi.org/10.1007/s10539-021-09831-0

De Wouters, d'Oplinter A. Huwart S. J. P. Cani P. D. & Everard A. (2022). Gut microbes and food reward: from the gut to the brain. *Frontiers in Neuroscience* 947240–947240. https://doi.org/10.3389/fnins.2022.947240

García-Cabrerizo R. Carbia C. O Riordan K. J. Schellekens H. & Cryan J. F. (2021). Microbiota-gut-brain axis as a regulator of reward processes. *Journal of Neurochemistry*1495–1524. https://doi.org/10.1111/jnc.15284

González-Arancibia, C. Urrutia-Piñones, J. Illanes-González, J. Martinez-Pinto, J. Sotomayor-Záratem R. Julio-Pieper M. & Bravo J. A. (2019). Do your gut microbes affect your brain dopamine? *Psychopharmacology* 1611–1622. https://doi.org/10.1007/s00213-019-05265-5

Han, W. Tellez, L. A. Perkins, M. H. Perez, I. O. Qu, T. Ferreira, J. Ferreira, T. L. Quinn, D. Liu, Z. & Gao, X B. (2018). A neural circuit for gut-induced reward. *Cell* 665–678.e23. https://doi.org/10.1016/j.cell.2018.08.049

Hassan S. (2013). *Freedom of mind: helping loved ones leave controlling people cults and beliefs* (First). Freedom of Mind Press.

Henrich, J, P. (2016). *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press.

Henrich, J. & Boyd, R. (2002). On modeling cognition and culture. *Journal of Cognitions and Culture*, 2, 87-112. https://doi.org/10.1163/156853702320281836

Henrich, J. & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology: Issues News and Reviews* 123–135. https://doi.org/10.1002/evan.10110

Heyes, C. (2016). Blackboxing: Social learning strategies and cultural evolution. *Philosophical Transactions of the Royal Society*, 371(1693), 1-8. https://doi.org/10.1098/rstb.2015.0369

Heyes, C. (2018a). *Cognitive gadgets: the cultural evolution of thinking*, Cambridge: Havard University Press.

Heyes, C. (2018b). Enquire within: Cultural evolution and cognitive science. *Philosophical Transactions of the Royal Society*, 373(1743), 1-9. https://doi.org/10.1098/rstb.2017.0051

Heyes, C. (2023). Rethinking Norm Psychology. *Perspectives on Psychological Science.*

Hong, Z, & Henrich, J. (2021). The cultural evolution of epistemic practices: The case of divination. *Human Nature, 32*, 622–651. https://doi.org/10.1007/s12110-021-09408-6

Joiner, J. Piva, M. Turrin, C. & Change, S. (2017). Social learning through prediction error in the brain. *Npj Science of Learning*. https://doi.org/10.1038/s41539-017-0009-2

Kendal J. Giraldeau L.-A. & Laland K. (2009). The evolution of social learning rules: payoff-biased and frequency-dependent biased transmission. *Journal of Theoretical Biology* 210–219. https://doi.org/10.1016/j.jtbi.2009.05.029

LaCroix, T. (2020). Evolutionary Explanations of Simple Communication: Signalling Games and Their Models. *J Gen Philos Sci* **51**, 19–43. https://doi.org/10.1007/s10838-019-09481-7

Laudan L. (1977). *Progress and its problems : towards a theory of scientific growth*. University

of California Press.

Lewens, T. (2015). *Cultural evolution: Conceptual challenges*. Oxford University Press.

Lewens, T. (2002). Technological innovation as an evolutionary process darwinnovation! *Studies in History and Philosophy of Science*, 195-203. https://doi.org/10.1016/S0039-3681(01)00038-3

Lewontin, R. C. (1970). The Units of Selection. *Annual Review of Ecology and Systematics*, *1*, 1–18. http://www.jstor.org/stable/2096764

Niv Y. Joel D. Meilijson I. & Ruppin E. (2002). Evolution of reinforcement learning in uncertain environments: a simple explanation for complex foraging behaviors. *Adaptive Behavior* 5–24. https://doi.org/10.1177/1059-712302-010001-01

Nichols, R., Moll, H., and Mackey, J. L. (2019). Rethinking cultural evolutionary psychology. *Journal of Cognition and Culture*. 19, 477-492 https://doi.org/10.1163/15685373-12340070

Norton, D. (2022). How analogy helped create the new science of thermodynamics. *Synthese*. 200 (4):1-42.

Novaes C. D. (2012). *Formal languages in logic: a philosophical and cognitive analysis*. Cambridge University Press. https://doi.org/10.1017/CBO9781139108010.

Popper K. R. (1979). *Objective knowledge: an evolutionary approach* (Rev.). Clarendon Press; Oxford University Press.

Radulescu A. Niv Y. & Ballard I. (2019). Holistic reinforcement learning: the role of structure and attention. *Trends in Cognitive Sciences*. *23: 278-292*. https://doi.org/10.1016/j.tics.2019.01.010

Railton, Peter. (2021). Normative guidance, evaluative Guidance, and Skill. *Analyse & Kritik*,

43(1), 235-252. https://doi.org/10.1515/auk-2021-0014

Richardson, J. P. & Boyd, R. (2004). *Not by Genes Alone: How Culture transformed Human Evolution*. The University of Chicago Press.

Sperber, D. (2000). An objection to the memetic approach to culture. In R. Aunger (ed.), *Darwinizing culture: The Status of Memetics as a Science*. (pp. 163-174). Oxford University Press.

Skyrms B. (2010). *Signals: evolution learning & information*. Oxford University Press.

Sterelny, K. (2012). *The Evolved Apprentice: How evolution made humans unique.* The MIT Press.

Sterelny, K. (2021). *The Pleistocene social contract: Culture and cooperation in human evolution*. Oxford University Press.

Strandwitz, P. 2018. Neurotransmitter Modulation by the Gut Microbiota. *Brain Research* 128–33. https://doi.org/10.1016/j.brainres.2018.03.015.

Sydnor, V. J. Larsen, B. Bassett, D. S. Alexander-Bloch, A. Fair, D. A. Liston, C. Mackey, A. P. Milham, M. P. Pines, A. & Roalf D. R. (2021). Neurodevelopment of the association cortices: patterns mechanisms and implications for psychopathology. *Neuron* 2820–2846. https://doi.org/10.1016/j.neuron.2021.06.016

Tomasello, M. (2014). A natural history of human thinking. Harvard University Press.

Tomasello M. (2019). *Becoming human: a theory of ontogeny*. Belknap Press of Harvard University Press.

Williams D. (2021). Socially adaptive belief. *Mind & Language* 333–354. https://doi.org/10.1111/mila.12294