doi:10.5281/zenodo.15581910 gavindavidyounglectures.au

The 15th Gavin David Young Lecture in Philosophy
The University of Adelaide, September 2024

Evolution and Moral Realism*

Take 2

Kim Sterelny†
Australian National University

Abstract

A decade ago I offered a defence of naturalistic moral realism against evolutionary debunking arguments. It had two elements. One was rejecting a dichotomy between evolutionary vindication and debunking in favour of a continuum. The second (a) proposed a metaethical principle on the basis of a hypothesis about the evolutionary drivers of moral thinking: a moral principle is true if it is part of a set of moral norms which, if acted upon, would optimise the cooperation profits of a community, and (b) argued that by this criterion, folk moral opinion was quite often approximately true. I still reject the dichotomy, but the rest of the argument was too friendly to realist naturalism in ethics. It mischaracterised the selective history of moral thinking, over-estimating its role as a cooperation amplifier, and under-estimating the epistemic challenges facing individual agents as they develop their moral opinions. The upshot is a shift towards the debunking end of the vindication-debunking continuum.

Keywords Moral realism · Evolutionary debunking arguments · Human evolution · Cooperation · Moral naturalism

*Thanks to the Department of Philosophy at the University of Adelaide for inviting this lecture and providing an occasion for a rethink, to the three (anonymous-to-me) referees for their constructive reports, and to the less anonymous David Copp for his extensive feedback.

† 0000-0003-3159-6698

1 Introduction

Recent decades have seen a lively debate on evolutionary error theories of moral thinking (for reviews, see Vavova 2015; Korman 2019; Levy and Levy 2020). These theories presuppose some version of cognitivism about normative claims. They see moral judgments as making truth-apt claims; judgements of right and wrong, fairness or unfairness are not disguised expressions of preference, instruction, or emotion. They seek to make objectively correct claims. But according to evolutionary error theorists, with the exception of logically trivial cases, these claims are never warranted.[1] Beliefs are warranted only if they are the result of epistemically reliable mechanisms, and for a mechanism to be reliable it must be counterfactually resilient.[2] Kate's belief that Toby (the dog) weights less than 25kg is warranted, for Toby is a canine member of the family. Had Toby weighed (clearly) more than 25kg, his human family, Kate included, would believe that Toby weighed more than 25kg. This counterfactual sensitivity is no mystery, as Kate's beliefs about Toby's weight are caused (partly) by his weight. Consider now an uncontroversial moral claim: it would be wrong to set Toby on fire for fun. Evolutionary error theorists think we have moral beliefs like this not because they are causal responses to independent moral facts, but because having this species of belief helps us successfully navigate our social environment. Especially in the small scale and interdependent communities in which our moral psychology evolved, we all had an interest in suppressing sadism (and almost all of us still have). Moreover, given others' attitudes to sadistic pleasures, most of us have had an interest in not being sadistic. As a consequence, moral beliefs of this kind were and are apt to increase our fitness. But that does not depend on the beliefs being true, or mostly true.

The fact that our belief-forming dispositions are the result of evolution is in itself no reason for scepticism about belief. Consider the fact that we are remarkably good at 'mind-reading': that is, at identifying the desires, intentions, emotions, and beliefs of others, especially those familiar to us. This skill is essential for smooth social interaction and it might well have a selective explanation. If so, that would be no reason to suppose that our beliefs about others' mental states are unjustified or mostly mistaken. Mind-reading enhances fitness because, and only because, it

[1] There is a variant: non-trivial claims are false, because moral facts cannot find a home in a naturalistic view of the universe, one in which we are evolved animals. A non-evolutionary version of this view is developed by Mackie (1977). An evolutionary version is in Ruse (1986).

[2] Here I rely on externalist views of knowledge: it suffices that the mechanisms that build belief are actually reliable, whether or not the agent in question can tell whether they are reliable.

usually enables us to correctly identify the thoughts of others. The mechanisms of that support mind-reading sometimes fail. We misread one another from time to time. But in familiar contexts, we interpret one another seamlessly, without even noticing this impressive cognitive skill. An evolutionary explanation of this cognitive capacity would be vindicating, not debunking. In general, it is plausible that belief-forming mechanisms are adaptive because they typically instal true beliefs, for action based on true belief is less likely to miscarry than action based on false belief. Cognitive mechanisms enhance fitness by enabling us to reliably track independent features of our environment, though of course given the finiteness of cognitive resources and the complexity of human environments, reliability is never perfect.

However, it is true that selection on cognitive capacities will not always result in truth-tracking mechanisms. For there are limits on a famous metaphor about belief: that they are the maps by which we steer. When belief guides practical interaction with our immediate environment, they are indeed the maps by which we steer, and map inaccuracy is apt to impose serious costs. One would expect selection to powerfully favour cognitive mechanisms that conduce to the accuracy of such beliefs. But especially with respect to those judgements that are the result of our immersion in the sea of public representations in which we swim, many of our beliefs do not guide action in ways that derail if those beliefs are false. Many millions of Americans believe (let's suppose falsely) that Trump was defrauded of victory in the 2020 Presidential Election. Very few, if any, embarked in a course of action which would have succeeded only if that belief were true. That is probably true even of those gathered on January 6 2021 to protest at and/or disrupt the formal recognition of the Biden victory. That attempt failed, and presumably it would have failed even if the election had actually been stolen. It failed because the key political institutions in Washington believed that Biden had won fairly, and would have failed even if that institutional belief was faulty. We have many beliefs – about the past, about our social environment, about the causal structure of our world – that either do not prompt action at all, or do not prompt actions whose success is truth-dependent. I believe that lithium is a light metal now important in battery manufacture, but since I am neither a battery engineer nor an investor in such technology, nothing I do will seriously miscarry if I am wrong. In the case of the political beliefs of the Trump supporters, these beliefs will mostly induce social signalling and other overtures to those likely to share their beliefs. These actions are quite likely to have positive consequences for the agent in question, building

supportive social relations with the like-minded, independently of the truth of the belief.

The Trump example illustrates this more general possibility: beliefs, when expressed, can systematically support positive social interactions independently of their truth. An important hypothesis about the evolution of religion illustrates this possibility, and it is an important model for the debunking view of moral judgement. Religious belief is not a human universal. But it is just about pan-cultural. Especially in small scale, pre-state communities, community life is characterised by religious belief systems typically paired with, and justifying, distinctive ritual practices. Jointly this religious package shapes the norms, customs and organisation of the community. The hypothesis is that these systems of belief and practice are typically adaptive for the community as a whole, and for most of the individuals in the community. The ritual life of the community contributes powerfully to local identity and cohesion, and hence promotes cooperative and prosocial interactions within that community. The belief system almost always commits agents to the existence of powerful occult agents who monitor and care about their actions. Since those agents respect and fear those occult beings, religious belief brings many of the benefits of efficient and impartial policing while paying few of the costs of such policing (see for example Sosis and Alcorta 2003; Sosis 2006; Whitehouse 2008; Whitehouse and Lanman 2014). On this picture, our ancestors were under selection for belief in supernatural agency independently of the existence of such agents. Selection tended to make theists of us; it would do so equally in a godless or a god-haunted world. This is a debunking argument. If successful, it does not show there are no gods. Rather, it shows that our readiness to believe is independent of theological phenomena. It is not true that: had the theological facts been different, our religious beliefs would have been different. Debunkers about moral norms offer a similar analysis of normative thinking.

Let's now turn to norms. I take these to be mutual expectations about choice and action, expectations that if violated result in sanctions of some kind, formal or informal. All known communities are regulated by social norms. It is less clear that moral norms are a cross-culturally identifiable and pan-cultural subclass of social norms. So some of the ideas under consideration here are specifically about moral norms; others are about social norms more generically. However, while evolutionary ideas about norms vary in important ways, most see them as contributing to our capacities for cooperation. They do so by facilitating high stakes cooperation (Tomasello 2016); by helping to stabilise cooperative interactions over

larger spatial or longer temporal scales (Joyce 2006; Kitcher 2011); by stabilising cooperation in larger social groups (Richerson and Boyd 1998; 1999; Seabright 2010; Bowles and Gintis 2011; Boyd 2018; Henrich 2018); and by signalling to others reliability as a social partner (Baumard, André, and Sperber 2013; Stanford 2018). Details vary with the specific model, but the overall shape is the same, and it parallels the ideas above about religion. By buying into objective moral truths, agents increase their fitness by enhancing their access to the benefits of cooperation. They tend to be trusted, or more trusted, in social interactions because they are actually more trustworthy. In environments in which information about others is freely available and quite reliable, the best way of seeming to be a good choice in cooperative interactions is to actually be a good choice. Real moral commitments and the emotions linked to those commitments make agents better able to resist the temptations to cheat. Likewise, they make those agents more willing to sanction those who do cheat, even when that cheating does not directly impact them. To be effective, these judgements must be sincere, and sincerely believed to track an objective fact. To resist the greedy temptation to make off with the whole kangaroo, you have to genuinely believe that only an utter shit would take it all. But the effect on temptation control does not depend on the judgement being true. If there were no moral facts, or different moral facts, it would still be fitness-enhancing to believe that there were such facts.

2 Moral Realism, Take 1

In a recent paper with Ben Fraser, I resisted the parallel between religious and moral beliefs. We argued that there was a pivotal difference between the explanations of religious and moral thinking (Sterelny and Fraser 2017). There is a credible hypothesis about the character of moral facts, and our cognitive response to such facts. If true, that hypothesis indicates that moral thinking enhanced fitness by truth tracking. The idea is simple. If moral thinking is an evolved response to the opportunities and challenges of cooperation in human communities, the moral facts derive from natural facts about community life. Some ways of interacting, and some practices are apt to stabilise and extend cooperation. Other practices are likely to limit and erode it. If moral cognition evolved to enhance our capacities to cooperate, we were under selection to identify, conform to,[3] and endorse

[3] To the extent permitted by temptation.

practices of the first kind. Likewise, we were under selection to identify, avoid, and disapprove of practices of the second kind. Fraser and I suggested that as a consequence, to a considerable degree, moral thinking satisfies the condition of counterfactual sensitivity. Practices that support or enhance cooperation do so in virtue of features our environments and of ourselves. If either of those had been different, cooperation-supporting practices would likewise have been different, and so too, Fraser and I supposed, would be our moral beliefs. If we had been very different – if we reproduced asexually – or if our environment had been very different – if we lived in social environments in which a successful life depended on hair-trigger willingness to inflict extreme violence – we would have had different moral beliefs.[4] To some significant degree, by natural design, our moral thinking tracks and responds to facts about cooperation.

Notoriously, moral opinions vary culturally. We saw this as a feature not a bug. For moral truths also vary culturally, if moral truth is determined by the feasible practices which, if conformed to, sustain cooperation. For such practices are sensitive to social and physical environmental parameters. For example, in some environments, the envelop of resources available to a community increases if individuals or families have incentives to invest in modifying the local environment: protecting fruiting trees by guarding them from animals, clearing weeds, making firebreaks, fertilising, or irrigating; building and maintaining stone fish-traps in streams and tidal estuaries. In those contexts, norms of private property, allowing privileged access to the products of these resource engines, might well be part of a cooperation-enhancing normative package. In other contexts, where variation in resource availability is driven by external factors under no-one's control, cooperation in the community might depend on norms of communal sharing; so-called 'demand sharing'.[5] Environmental variation in best practice was important, and so our ancestors were under selection to track and respond to cooperation-enhancing practices, whatever those practices were or are. Selection did not favour recognising, responding to, and genetically wiring-in a specific set of Late Pleistocene customs.[6]

4 There are social environments which somewhat approximate this condition: environments with no or little state control, and where resources crucial for survival are mobile and easily seized: herds are the prime example. In these environments, deterrence – the threat of lethal violence – must be credible.

5 It remains possible of course that very general principles characterise all these varying packages of cooperation-supporting practices; perhaps principles of reciprocation, or family support, are always present in some form.

6 Richard Joyce (2006) argued against evolutionary naturalism on the grounds that it would make moral truths truths about Pleistocene social arrangements. On the view defended in our paper, the

We did not claim that all variation in moral opinion reflected local variation in moral facts. But we did suggest that this factor was part of the explanation of normative difference. Moreover, normative variation helps explain how agents can identify practices that promote or erode cooperation. Connections to other communities through intermarriage and other links allow information about how other communities live to filter into common knowledge. Both our recent and more remote ancestors were able to observe variation in social practices, and so had access to information about customs and habits of interaction with others that work well, and badly. Of course, truth tracking comes in degrees, and in two ways. Practices themselves can be more or less ideal in sustaining cooperation and maximising its benefits. Ideally, perhaps, norms of sharing might require enough sharing to manage risk and to reduce conflict while still incentivising individuals to invest in productive effort. A less ideal principle – share large food packages, but not small ones – might satisfice, helping to manage risk and reduce social tensions. Agents might do reasonably well at identifying and committing to practices that satisfice in their local conditions, but find it very difficult to identify ideal ones. Without committing to any specific threshold, Fraser and I suggested reasonable competence in identifying satisficing practices. There will be noise, contingencies of local history, and bias. But to some degree, when moral facts are different, moral opinions were different too.

Fraser and I did not claim to have vindicated evolutionary naturalism. Rather, we argued that the contrast between vindicating and debunking genealogies should be reconceptualised as a continuum. Systems of folk opinion can be partially vindicated mosaics of error and successful tracking. The earth-centred astronomy of the ancient world, which we took to be a regimented formulation of folk belief about the heavens, was our model. Much of the theoretical apparatus of Ptolemaic astronomy was just wrong. But foundational concepts corresponded to real differences in the heavens, and the observational dynamics were represented accurately enough to support navigation, the anticipation of tidal movements, and seasonal changes. If the moon's orbit around the earth had taken 40 days to complete, beliefs about tidal periodicity would have varied appropriately. Without being fully vindicated, Greek astronomy is much better credentialed than Greek theology. We suggested a similar assessment of folk moral thought. The general

moral truths were based on practices that would support cooperation in the context of the actions being assessed. On that view, there are moral truths about life in contemporary environments.

background theory is often dodgy, as in many communities, norms are given their public justification by being embedded in mytho-poetic narratives. But we suggested that the normative packages of many communities seemed often to be reasonably well tuned to enhance cooperation in the face of potential destabilisers; an eyeball reading of the ethnography that was supported by the more systematic data of Curry and colleagues, who showed that many communities, perhaps all communities, endorse plausibly cooperation-promoting social norms (Curry, Mullins, and Whitehouse 2019).

There is a live tradition in philosophy that suggests, on general metaphysical grounds, that no form of reductive naturalism could be true (Horgan and Timmons 1991). While deeply sceptical of such metaphysical arguments, nothing in the paper with Fraser addressed this line of thought. But with respect to the specifically evolutionary rejection of naturalist realism, our Take 1 depended on a conjunction of three ideas:

1. In a given physical, biological and social environment,[7] there will be feasible packages of social practices which, if adopted, would stabilise and extend inter-agent cooperation.
2. Capacities to recognise and commit to moral norms (or more generally, social norms) have evolved to help agents tune their decision making so they can enjoy the profits of honest cooperation. Moral thinking evolved to help us reap the benefits of cooperation, and so, moral opinions at a time and place are true if they endorse cooperation stabilising-practices, or condemn cooperation undermining ones. These need not be ideal. We can count as true the opinion that one should share the meat of a deer, but one can keep a rabbit for one's own family, even if a more nuanced sharing norm would do even better.
3. Reliability comes in degrees. While it is clear that our moral opinions do not perfectly track the local moral truths, as identified above, we thought that in many communities, humans track them fairly well. After all, in almost all human communities, social life continues to be remarkably cooperative.

We took this to be a somewhat vindicating package. To some significant degree, our moral opinions are reliably true.

[7] Some features of a social environment are not readily alterable: its economic base; the skill set and embodied capital more generally of those within it; community size and connectedness; age profile. Since these features remain stable (catastrophe aside) over years and decades, it is reasonable to treat them as fixed contexts of action.

3 Retreat: Revising the Selective Hypothesis

I have become much more dubious about the second and third elements of this trifecta, and I will begin with second thoughts on moral facts. The selective hypothesis on which Fraser and I depended needs serious revision. Our view of moral realism embraces a version of metaethical naturalism: moral facts *just are* facts about cooperation. This proposal rested on an evolutionary hypothesis about the selective history of moral thinking together with a more general view about the relationship between mental representation and biological function. In general, the content of a belief is fixed by its evolved function. Our capacity to form moral beliefs is the result of selection to identify and respond to facts about cooperation, so those beliefs are about cooperation-facts. We thus bought into some generic version of teleosemantics, a set of views about the ways selective history determines representational content, a framework developed in most detail by Ruth Millikan (1984; 1989) and Nick Shea (2018). But we took the view that moral opinions are opinions about cooperation to be supported by a second consideration: the proposal was reasonably consistent with folk moral thinking.

Our reasonable-consistency claim involved an application of the 'Canberra Plan' (Jackson 1998). One important project within the Canberra Plan was to explore the status of domains of folk thought in the light of our best science. It comes with a method for evaluating the relationship between a domain of folk thinking and our best scientifically informed model of the world. Begin with a regimentation of the target domain of folk opinion by identifying the most important and least contentious folk opinions in that domain, the 'platitudes'. The folk zoology of animals might include such platitudes as: animals are alive; they can move autonomously; animals need to eat to live; animals come in many varieties but the varieties do not blend; an animal comes into existence only through being born from one of its own kind; and so on. We then ask, from the perspective of our best science, whether anything in the world satisfies these core folk opinions. If there is, a domain of folk opinion is vindicated, as folk zoology probably is. If there is not, as in the contrasting case of 19th century European racial theories, it is debunked. We should probably embrace an error theory of such racial classification. There are no races, as those Europeans thought of races.

Importantly, there are intermediate possibilities, and on our analysis, moral opinion was one. In most communities (perhaps all), core moral opinion endorses important modes of cooperation; for example, helping those who have helped you;

or being generous to unlucky members of your community. But it is also true that just about all moral codes include moral precepts with nothing or less than nothing to do with cooperation; for example, puritanical norms of various kinds. Allen Buchanan has pressed this point, arguing treating moral thinking as if it were just about cooperation ignores too much of moral opinion (Buchanan 2020). Moreover, many communities endorse far from cooperative practices towards outgroups, minorities and underclasses. In the light of such cases, Fraser and I accepted that in many, perhaps all communities, there was at best a rough match between their central moral opinions and the moral facts as we identified them. So the Canberra Plan case for identifying the moral facts with cooperation facts delivers only a partial match. So moral thinking originated through selection favouring the capacity to track and act on cooperation facts. Some, perhaps much, but not all moral thinking continues to be about those facts,

I have subsequently come to see that the case for identifying moral facts as cooperation facts is eroded by a more nuanced view of the selective history of moral thinking. Selection for tracking and responding to cooperation-enhancing practices turns out to be only one of the drivers of the evolution of moral thinking. Norms, and agents' commitment to norms, function as coordination devices and social signalling mechanisms. Human lives are interdependent, making it important to coordinate with other agents. In intimate social circles, this can be based on common knowledge. Each knows the others' habits and preferences. As social scale and complexity increases, this becomes a less feasible basis for mutual anticipation. Agents rely more on the formal and informal norms regulating social behaviour. There is less friction and misfire in social interactions if agents have similar expectations about such matters as: the organisation of work; the spatial organisation of campsites or villages; how collective decisions are made; how stored resources are distributed; how children become full members of the community, and so forth. As Robert Boyd has pointed out, once the social world becomes at all complex, without appropriate social tools, conflicts can easily flare even amongst agents who take themselves to be acting prosocially. For what counts as being prosocial? How much of a valued resource should be shared and with whom? Who is responsible for tending a fire that escaped control? Such issues are inherently ambiguous. Norms reduce such ambiguity substantially. They enable agents to know what is expected of them, and what they can expect of others (Boyd 2018). More generally, they make the social world more predictable. Norms reduce

uncertainties about others' actions, making every form of planned action more reliable.

Norms that facilitate coordination by making social interaction more patterned and predictable of course also facilitate cooperation, for much cooperation requires coordination. But many forms of coordinated social interaction have nothing to do with cooperation; if there is a norm of meeting in the kitchen for morning tea at 11am, and I am in the mood for a quiet morning, I know when not to go. Human lives are so connected that the reduction of social ambiguity is relevant to just about everything we do. Moreover, Cailin O'Connor has argued (on the basis of extensive modelling) that in an important class of cases coordinating norms reduce cooperation incentives by entrenching unfair divisions of the social surplus (O'Connor 2019). When communities are divided into groups (by age, gender, social role, occupation, whatever); when a mode of social interaction is predicated on group membership; and when groups are of unequal size, there is a tendency for smaller groups to be exploited by larger ones. Exploitation tends to emerge because it is more expensive for members of the smaller group to opt out of interacting with members of the larger group than vice versa. All else equal, their bargaining position is weaker, and they tend to end with a smaller slice of the pie. Once this pattern establishes, it readily acquires normative status, as agents predicate their own choices on expectation of conformity to that pattern, and impose social costs if unpleasantly surprised. The upshot is that the efficient solution of coordination problems is one selective driver of the evolution of our capacity to notice, endorse and conform to social/moral norms. Norms that reduce social ambiguity and which stabilise as solutions to coordination problems need not enhance cooperation.

So normative thinking is not just a cooperation-enhancing adaptation. A second complication emerges from two ways of meeting challenges to cooperation. In many circumstances, cooperative social customs are liable to erode if and as agents succumb to temptations to defect, enjoying the benefits of cooperation without paying its costs. Models of cooperation see this challenge as being solved in two different ways, with normative commitments playing somewhat different roles (though most real cases are mixed). Partner control models take cooperative patterns to be stabilised by agents incentivising their social partners to cooperate, and/or by reducing their incentives to cheat. On this family of views, exemplified in the analyses of Richerson and Boyd (2001; Richerson, Boyd, and Henrich 2003), moral norms play important roles in defining what counts as cheating, and mobilising and motivating coalitional, hence inexpensive, punishment of cheats.

Self-control is part of this same package, in part induced by prudential avoidance of anticipated punishment, in part intrinsically motivated by subjective commitment to the appropriateness of these norms. In contrast, partner choice models take cooperation to by stabilised by assortative interaction (for a clear introduction to this way of thinking, see Witteveen 2021). Willing co-operators recognise and preferentially interact with others, leaving would-be cheats with the dregs. On this picture, developed in some detail by Nicholas Baumard with various collaborators, moral commitments are signals (Baumard, André, and Sperber 2013; Fitouchi, André, and Baumard 2023). He argues that if this partner choice picture is right, the most critical moral norm will be commitment to fair dealing, and that there is empirical evidence (mostly from experimental economics) that this norm is indeed central. These two perspectives on the relation of normative commitment to cooperation both fit the more general thesis that moral thinking has evolved to enable us to identify and grasp opportunities for mutually profitable cooperation. On partner control models, this works by incentivising social partners to adopt cooperation-supporting social practices. On partner choice models, it works by helping us to recognise others who are antecedently willing to adopt such practices, by advertising our own commitment to them.

However, public commitment to norms does not just signal commitment to cooperative social practices. These are often instead signals of identity, of membership of a local group (for an introduction to the role of such signals, see Fessler and Quintelier 2013). An Arsenal supporter signals commitment: by wearing appropriate insignias to the game; by travelling and sitting with other supporters; by loud enthusiasm for team success, and in some cases by expressing and acting on norms enjoining verbal and physical conflict with supporters of rival teams. This social mechanism is ancient. Humans do not just live in communities that operate as relatively cohesive social and economic units. Individual agents within those communities are aware of their membership of these communities, value this membership, and see their communities as distinct from others. Especially in small world pre-state societies, this collective identity is both strengthened and advertised by a shared ritual life; a shared set of origin stories explaining both their specific identity and rights to place; often a distinctive language or dialect; distinctive material symbols, expressed in styles of dress, markings of place, personal adornment (Whitehouse 2021).

The archaeological record of material symbols suggests that this mode of expressing and reinforcing local identity has been an aspect of human life for

100,000 years or more, and it could be much older. For as Stibbard-Hawkes (2025) has pointed out, the material signs of this aspect of our social lives are particularly likely to disappear from the physical record with the passage of time. Norms play important roles in establishing and signalling community identity and difference. Communities vary in their food taboos; norms of family structure (whether the husband joins the wife's family or vice versa); norms of appropriate social behaviour (like mortuary norms); norms governing sexual behaviour and marriage. Many Australian Aboriginal communities, for example, have complex systems in which every individual in the community is born into a specific section or subsection, and can marry only into a designated other (sub)section (for a relatively simple example, see Gould 1969). Public conformity to these norms, and to other public signals of identity, are essential to good repute within one's home community. That is always important, and is essential in social environments in which inter-community interactions are fraught. In those social worlds, without the backing of your community, your prospects are dire. These signals of commitment and identity, by promoting trust within the community of one another, may well function to enhance local cooperation. But the content of these norms, and hence the signal given by public commitment to and conformity with them, often have nothing to do with cooperation. The norm might be one that mandates a particular organisation of camp-sites, or the physical punishment of children.

If norms have the role of defining and signalling local identity, there is no expectation that such norms will track cooperation-promoting practices. Indeed, it is arguable that selection might favour signal divergence that push local norms of the kind that signal identity away from optimal practices. If the folk across the river happen to have hit on well-functioning marriage rules (perhaps would-be grooms perform some but not too much bride service), and there are factors encouraging your community to emphasise its difference from those slatterns and never-do-wells across the waters, there will be some tendency to adopt different marriage rules. Once established, these norms are difficult to change, as non-conformity attracts social costs. Be that as it may: the main point is that to the extent that normative commitments serve as social identity markers, they will not systematically track good cooperative practices, for identity markers can be, and often have been, arbitrary: dietary taboos; the organisation and design of villages; norms of dress and appearance. It follows that if selection has shaped our minds to notice and adopt norms of our local community through the benefits of more seamless coordination

and local trust, we have not thereby been under selection to notice and adopt practices that promote cooperation.

The upshot of this section is that the case for identifying moral facts with natural facts about the practices that promote or undermine mutually beneficial cooperation is less strong than I had supposed. A more nuanced consideration of the selective factors suggests that tracking cooperative practices was only one factor selecting for normative thinking. Given the importance of both local coordination and signalling and committing to local identity, it may not even be the most important factor.

4 Retreat: Can the Folk Recognise the Moral Truths?

A defence of naturalistic moral realism requires two elements. One is a credible account of the character of moral facts, showing them to be a species of natural fact. A second is a credible moral epistemology, showing that we have decent prospects of recognizing those facts, of acquiring moral knowledge. Fraser and I were not naïve about our ancestors' abilities to identify practices that promoted (or eroded) cooperation. Social environments were not transparent, and agents in those environments had limited resources for devoting their attention to how their communities were run. So actual moral opinions would never perfectly, or even approximately, match local moral truths. Nevertheless, we thought there would be sufficient local variation, and sufficient information from neighbouring communities, for there to be a noisy but real tendency for moral opinions to reflect truths about good local practices. While I still think there is such a tendency, it is significantly weaker than I thought. Most of our evolutionary history has played out in small mobile forager communities. In such communities, little information is available that would enable a community to tune its normative lore. Signal is obscured by noise. In many parts of the world, the Holocene saw a shift to more sedentary, larger and less equal social worlds. In such communities, the informational environment was biased: those playing a pivotal role in transmitting lore to the next generation were positioned to bias normative education in ways that favoured their own interests.

Let's begin with noise. Social engineering is hard, for all human communities are complex, and the effects of norms and customs interact with one another and other aspects of community life, and over long-time delays (see for example Gauss 2021). Social environments are causally opaque. The physical and biological environment

is informationally challenging too, but here trial and error experimentation can be a powerful tool, especially when trials are inexpensive, making multiple trials possible. Likewise trial-and-error is epistemically powerful when effects, positive and negative, have short time delays and appreciable magnitudes, so they are salient. Broadcast matters too: information flows best when both the experiment and its effects are public, so available for third-party observation. In her work on a Madagascan fishing community, Rita Astuti describes a paradigm of this engine of local adaptation: a single individual introduces a new configuration for rigging the sails and masts on his boat, allowing more efficient tacking, and hence better control of direction. Within about a decade, the innovation went to fixation (Astuti 2024). No surprise. The new configuration did not require a new boat design or different fishing equipment: the configuration could be altered independently of the rest of the fishing rig. The effect of the innovation was public, repeatable, and rapid. No social costs were paid in switching to the new configuration; while fishing was central to community identity, sail configuration was not. There was no collective action problem: each fisher could switch if and when he so chose. There were not material costs in switching: the existing gear was re-purposed. No new skills were needed, so there was no efficiency or productivity trench between abandoning the old configuration and becoming competent with the new one; as there would be, for example, in switching from javelin hunting to bow hunting.

None of this is likely to be true of norms. Deliberate innovation is essentially impossible: that would just be seen by others as norm violation. As an individual, you can trial a new fishhook. But you cannot trial a new marriage rule. There are natural experiments: the differing practices of adjacent local communities (or of subgroups nested in the home community). But except for truly catastrophic practices, these will not be very informative. The effects of these variations on social life are likely to be felt over long time-frames, with little obvious connection between cause and effect, especially as it is unlikely that adjacent groups will vary normatively in only a single respect: in, say, their marriage rules. Norms are often legitimised from being nested in mythic narratives of origin, identity and right to place. Specific norms around marriage are often parts of linked systems of belief about family, gender and kinship. Where there is variation is one of these aspects of social regulation, there will typically be variation in others. The same facts about how norms acquire their legitimacy make normative change a very difficult collective action problem. If it is true that, for instance, the norms about marriage are seen as important to local identity, and are seen as important in demarcating that identity from others, the

social costs in trying to change them are likely to be high and the chances of success low.

While thinking that for the most part norms promote prosocial behaviour, in their work on our social life, Boyd, Richerson and Henrich place little reliance on individual capacity to adaptively tune social norms. Moreover, they note that evolutionary modelling suggests that a wide variety of social strategies can stabilise as Nash Equilibria, including markedly maladaptive ones (in part, because punishment can stabilise just about any norm). Yet they think markedly maladaptive practices are relatively rare, and prosocial ones relatively common. How so? They conceptualise this as an 'equilibrium selection problem' and suggest that cultural group selection has favoured groups that have lucked on norms that encourage prosocial interactions, and penalised groups whose norms encourage every-agent-for-themselves practices (Richerson and Boyd 2001; Henrich 2004; 2006). Evolutionary biologists have typically been sceptical of group level explanations of cooperation, as they expect cooperation in cooperative groups to be eroded by individual level selection for grasping the benefits of cooperation while dodging its costs (the classic exposition of this scepticism is Williams 1966). Boyd, Richerson and Henrich point out that this reason for scepticism does not apply to their model. For it envisages selection amongst groups all of whom are internally stable; no individual within any group can benefit by changing social strategy. It follows that no individual level selection counteracts the group level process by penalising individuals behaving prosocially and rewarding those acting anti-socially. There is no tendency for cooperation in groups already cooperative to decay.

It is true that if local communities are at equilibrium (a substantial 'if'), the effects of community level selection for prosocial norms are not undercut by individual selection for defection. Even so, selection on communities is unlikely to be a powerful, optimising force. Selection is effective when traits vary independently of one another (else you cannot fine-tune, say, child-rearing norms independently of gender division of labour norms); when the population is large, for chance has large effects on small populations; when heritability is high, and when the selective environment is stable over many generations. Without a stable selective environment, favourable traits at a time will not go to fixation, and nor is the incremental evolution of adaptation possible (for a nuanced discussion of the factors making populations strongly or weakly responsive to selection, see Godfrey-Smith 2009). Selection at the level of communities does not satisfy these conditions. In pre-state social worlds, most communities interact only with a few neighbours.

Moreover, relative to individual life spans, communities are quite long-lived, so generational turn-over, with extinction and founding new communities by colonization is relatively infrequent, while relative to the pace of community generational turnover, both social and physical environments are quite changeable. Relations with other communities, for example, might transition rapidly from being cooperative to quite hostile or vice versa, and that would have implications for local economic practices, for out-group marriage, and for norms about violence and conflict. Selection on the population of communities might well filter out those with grossly dysfunctional norms, but not much more.

We noted above that a convincing defence of naturalistic moral realism needs to include a credible moral epistemology, showing that we have decent prospects of acquiring moral knowledge. This section has suggested that for much of our evolutionary history, our history of life in small, pre-state communities, if moral facts are facts about cooperation enhancing or inhibiting practices, epistemic access to moral facts was decidedly problematic. This conclusion should not be over-stated. We are astute social animals, and our experience of life teaches us a good deal about what angers and dismays others, and what pleases them. That said, local custom plays an immersive role, with humans habituating and treating as normal wildly different social patterns. I doubt that there are any communities in which lethal violence is regarded as no more than boyish high spirits. But consider the enormous variation in community attitudes towards non-lethal violence: violence to and amongst children; of men towards women; between men and between women. Think too of the enormous variation in the extent to which community norms support or suppress female sexual choice. These varying practices are all normalised within the communities they characterise.

All human communities depend on cooperation, and so all have avoided norms that would cause catastrophic cooperation failures. In all communities there is some ability to recognize and support cooperation sustaining practices. Within that broad constraint, normative diversity is extensive, and while some is doubtless tuned to local environmental conditions, much probably is not. It is not plausible, for example, that varying norms about the physical punishment of children (from the regularly required to the utterly forbidden) reflect variation in local conditions. Likewise, norms about whether adolescent girls can choose their own sexual partners; allowed by the norms of the foragers of semi-arid lands in Africa, but not the semi-arid foragers of Australia. It is very easy to overstate human ability to recognize and correct dysfunctional norms. To some degree, our

inability to recognise dysfunction flows from the fact that the effects of norms are often long downstream and hard to recognize. Think, for example, on how difficult it has been to determine the effects of the physical punishment of children even in contemporary societies, with their ability to store data, keep long longitudinal records, and to analyse immense data sets with sophisticated statistical tools. But it is also due to our great capacity to adapt to and normalize very varying experiences, treating social practices as just the way things are.

Bottom line: in small scale pre-state societies, our ancestors had constrained and noisy access to moral facts (if those facts are about locally good or optimal cooperation practices). Their prospects for moral knowledge were limited. Social worlds are complex, and the effects of specific practices are difficult to estimate, even if all relevant parties are disinterested. But they are not. Sometimes agents with special influence over the normative beliefs of their community have interests of their own, interests that run counter to those of others. This adds bias to noise.

5 Retreat: Adding Bias to Noise

In many communities, norms include those of deference and respect for those in positions of authority, and of obedience to more formalised and explicit legal and quasi-legal codes. In policed societies, it is typically prudent to conform to those codes, but in such societies the institutions of normative education moralize these codes and obedience to them. These include: priests and organised religion; Elders, schools where they exist, centrally organised public rituals, commemorations and displays. Conformity is a moral or social duty, not merely prudential avoidance of retaliation. Typically, these norms of deference and obedience serve the interests of local elites rather supporting the distribution of the social surplus in ways that incentivise those in the community to invest time, energy, skill and resources in productive cooperation. Once states and their attendant institutions develop, with their control over the structured dissemination of lore, and their influence on the public representational space of the community, elite influence over the normative opinions of the community is to be expected. But systematic sources of bias pre-date states and their powers. Even in the relatively egalitarian social worlds of Australian Aborigines, transmission of lore was largely oblique, funnelled through high-prestige initiated males, Elders, and it is no surprise that the lore so

transmitted leveraged special sexual privileges for Elders (for a detailed exploration of the effects of these norms, see Keen 2006).

Humans are not cognitively autonomous. In developing and deploying our cognitive capacities we depend on adapted learning niches, organised and supported social learning, and especially in recent millennia, material aids to thinking (this claim is defended in detail in Sterelny 2012; 2021a). We have scaffolded minds. This has a momentous upside, making possible language and literacy, precise mathematical reasoning, and giving us access to information beyond the reach of our perceptual experience. But when scaffolding is in the hands of others, and those others have divergent interests, that scaffolding opens up possibilities of systematic manipulation (Timms and Spurrett 2023). Normative education is a site of such manipulation. The shift to more sedentary, larger, less equal social worlds is coeval with a shift to more alienated modes of scaffolded development. Children in mobile forager communities have a great deal of autonomy. They do not choose what they learn: they need to, and do, acquire the skills and lore of their community. But they have a good deal of control over how, when and with whom they learn. Much of their childhood is spent in mixed age play groups, with much learning by experiment, often collaboratively, and with peer-to-peer learning, all at their own pace. This development is supported by adults through the provision of material supports, advice, modelling, and occasional explicit teaching. Children in subsistence farming communities have lost much of this autonomy, and as learning and teaching becomes more institutionalised in larger mass societies, children have even less (see Sterelny 2021b for a review of the ethnographic support for these claims). In formal educational institutions, children have very little say about what they learn, when, and from whom. As the modal human experience has become one of life in a larger social world, there have been two changes in rough synchrony. Social worlds have become larger, more complex and less equal, and children and their families have had reduced control over the informational scaffolding of developmental trajectories to adulthood. This development allowed the malign alliance of priests and chiefs to bias normative education in their own interests.

No doubt the actual powers of elites to shape moral education varied enormously over time and place. There is occasional evidence, from inquisition record, that opinion on the ground could vary very dramatically from official orthodoxy (see for example Ginzberg 2013). But where elite-controlled institutions could largely shape what could safely be said in public, they inserted not noise but bias into the factors shaping moral opinion. Of course, agents are not helpless in the face of

manipulative scaffolding. Dan Sperber and Hugo Mercier (especially) have shown that we have metarepresentational mechanisms of epistemic vigilance. These allow us to explicitly represent and assess incoming public opinion. The great demographic transition in Western Europe of the 19th and 20th centuries, with a great fall in family size, took place despite the natalist propaganda of church and state. We do not believe everything we are told, and this constrains manipulative options (Sperber 2001; Sperber, Clément, *et al.* 2010; Mercier and Sperber 2017). As DeScioli and Kurzban point out in their account of the evolution of moral thinking as a social support recruitment tool, bias typically has to be disguised in some way (2013; see also DeScioli 2016). Norms must be presented as if they were impartial and to the benefit of all; or at least as having an authoritative source independent of those who might benefit from a norm. These are the deliverances of ancestors or similar unquestionable authority figures. It is not the king but God himself telling us all to obey the king.

However, while epistemic vigilance is important its scope is limited. Metarepresentation and assessment requires the limited resources of attention and focus. Moreover, this cognitive apparatus is at best rudimentary in the young, and these are targets of moral education. While those young children can revise their initial moral opinions, that too is expensive. In his discussion of the social transmission of religious practices, Joseph Henrich notes that adult credibility is enhanced when adults engage in costly practices that only make sense if they genuinely believed the narratives they transmit – he calls these credibility-enhancing displays (Henrich 2009). An adult showing deference to rank, or sharing their resources with others, or respecting norms of social avoidance, all display to the children they teach their honest adherence to these norms. Moreover, targeted scepticism is difficult because most of the sources of noise discussed in the previous chapter remain important in larger social contexts: aside from obviously catastrophic options, it is hard to recognise the cooperation truths.

On some views, in recent human history, there has been genuine moral progress: for instance, outlawing chattel slavery (Buchanan and Powell 2018). If so, perhaps that shows improving access to moral facts. It is true that in the more recent past, we have developed new epistemic tools. We have permanent, curated, publicly accessible large data sets describing the state of community populations over time; we have statistical tools that allow us to estimate the consequences of social changes of many kinds; we have institutions which make it possible to collectively advocate for normative change. We have seen this in recent decades with community-wide

changes in normative views about same sex marriage and meat-eating. We have normative specialists: public intellectuals and others who devote most of their time to identifying maladaptive practices and advocating for normative change. However, these new epistemic tools have come on stream at the same time that societies have become larger, more interconnected and more complex. When many causal pathways interconnect, it is very difficult to predict the medium-and-long term consequences downstream from a point of origin. Consider a pertinent current example. 'WEIRD'[8] countries may be on the cusp of a normatively driven change from a past in which sex and its associated social roles came in just two flavours (with rare exceptions regarded as medical pathologies), and fixed at birth, to a future in which the equivalent of sex is not fixed at birth, coming in many more varieties, and is changeable. If WEIRD countries do change this way, it will be a change largely driven by normative arguments about individual rights to self-identify. The medium to long term consequences of this transformation are unknown, and in the view of some, unknowable. Gerry Gauss, for example, has recently defended the view that in large complex social worlds, the consequences of significant change are inscrutable (Gauss 2021). If that is even approximately right, the additional leverage gained from new epistemic tools is largely cancelled out by the increased difficulty of projecting change through complex sensitively interlocking systems.

This idea should not be overplayed: it is hard to see how the social consequences of alleviating dire Aboriginal social inequality could be bad in the long term, or how a savage reduction in health care and life expectancy of poor regional Australians could be good in the long run. We can have reasonable confidence in coarse-grained estimates of social consequences in limiting cases, especially when there have been similar phenomena in the past whose consequences we have noted, good and bad. But perhaps not much more.

6 Conclusion

Folk opinion about aspects of their environment can be largely right, hopelessly wrong, or somewhere in between. Despite the many retreats in this paper, I still locate folk normative opinion as somewhere in between, though now closer to the debunking end of the spectrum from debunking to vindication. 'It is wrong to set fire to dogs for fun' is in better epistemic shape than 'Jill is an octoroon' or 'James

[8] '*W*estern *E*ducated *I*ndustrial *R*ich *D*emocratic'.

is a mulatto'. But it is not in great epistemic shape. Realism requires a defence of the claim that a domain is characterised by objective mind-independent truths, and that we have some epistemic access to these truths: fallible, but somewhat reliable, perhaps improving. The prospects for both claims have dimmed. The claim that the evolved function of moral thinking is to represent cooperation-promoting practices turns out to be much over-simplified. Moreover, in all or most of the environments of human experience, outside extreme cases (and dog-torture may be one) our ability to recognise such practices is much more limited than we are apt to suppose.

References

Astuti, Rita (2024) 'Innovation and Social Identity in Madagascar', in Mathieu Charbonneau, ed., *The Evolution of Techniques*: 27–36. MIT Press.

Baumard, Nicolas, Jean-Baptiste André, and Dan Sperber (2013) 'A Mutualistic Approach to Morality: The Evolution of Fairness by Partner Choice', *Behavioral and Brain Sciences* **36**: 59–78. doi:10.1017/S0140525X11002202.

Bowles, Samuel and Herbert Gintis (2011) *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton University Press.

Boyd, Robert (2018) *A Different Kind of Animal: How Culture Transformed Our Species*. Princeton University Press.

Buchanan, Allen (2020) *Our Moral Fate: Evolution and the Escape from Tribalism*. MIT Press.

Buchanan, Allen and Russell Powell (2018) *The Evolution of Moral Progress: A Biocultural Theory*. Oxford University Press.

Curry, Oliver Scott, Daniel Austin Mullins, and Harvey Whitehouse (2019) 'Is It Good to Cooperate? Testing the Theory of Morality-as-Cooperation in 60 Societies', *Current Anthropology* **60**: 47–69. doi:10.1086/701478.

DeScioli, Peter (2016) 'The Side-Taking Hypothesis for Moral Judgment', *Current Opinion in Psychology* **7**: 23–27. doi:10.1016/j.copsyc.2015.07.002.

DeScioli, Peter and Robert Kurzban (2013) 'A Solution to the Mysteries of Morality', *Psychological Bulletin* **139**: 477–96. doi:10.1037/a0029065.

Fessler, Daniel M T and Katinka Quintelier (2013) 'Suicide Bombers, Weddings, and Prison Tattoos: An Evolutionary Perspective on Subjective Commitment and Objective Commitment', in Brett Calcott, Ben Fraser, Richard Joyce and Kim Sterelny, eds., *Cooperation and Its Evolution*: 459–83. MIT Press.

Fitouchi, Léo, Jean-Baptiste André, and Nicolas Baumard (2023) 'Moral Disciplining: The Cognitive and Evolutionary Foundations of Puritanical Morality', *Behavioral and Brain Sciences* **46**: e293. doi:10.1017/S0140525X22002047.

Gauss, Gerald F (2021) *The Open Society and Its Complexities*. Oxford University Press.

Ginzberg, Carlo (2013) *The Cheese and the Worms: The Cosmos of a Sixteenth-Century Miller*, 2nd edition. Johns Hopkins University Press.

Godfrey-Smith, Peter (2009) *Darwinian Populations and Natural Selection*. Oxford University Press.

Gould, Richard A (1969) *Yiwara: Foragers of the Australian Desert*. Collins.

Henrich, Joseph (2004) 'Cultural Group Selection, Coevolutionary Processes and Large-Scale Cooperation', *Journal of Economic Behavior & Organization* **53**: 3–35. doi:10.1016/S0167-2681(03)00094-5.

Henrich, Joseph (2006) 'Cooperation, Punishment, and the Evolution of Human Institutions', *Science* **312**: 60–61. doi:10.1126/science.1126398.

Henrich, Joseph (2009) 'The Evolution of Costly Displays, Cooperation and Religion: Credibility Enhancing Displays and Their Implications for Cultural Evolution', *Evolution and Human Behavior* **30**: 244–60. doi:10.1016/j.evolhumbehav.2009.03.005.

Henrich, Joseph (2018) *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press.

Horgan, Terence and Mark Timmons (1991) 'New Wave Moral Realism Meets Moral Twin Earth', *Journal of Philosophical Research* **16**: 447–65. doi:10.5840/jpr_1991_19.

Jackson, Frank (1998) *From Metaphysics to Ethics*. Oxford University Press.

Joyce, Richard (2006) *The Evolution of Morality*. MIT Press.

Keen, Ian (2006) 'Constraints on the Development of Enduring Inequalities in Late Holocene Australia', *Current Anthropology* **47**: 7–38. doi:10.1086/497672.

Kitcher, Philip (2011) *The Ethical Project*. Harvard University Press.

Korman, Daniel Z (2019) 'Debunking Arguments', *Philosophy Compass* **14**: e12638. doi:10.1111/phc3.12638.

Levy, Arnon and Yair Levy (2020) 'Evolutionary Debunking Arguments Meet Evolutionary Science', *Philosophy and Phenomenological Research* **100**: 491–509. doi:10.1111/phpr.12554.

Mackie, J L (1977) *Ethics: Inventing Right and Wrong*. Penguin Books.

Mercier, Hugo and Dan Sperber (2017) *The Enigma of Reason: A New Theory of Human Understanding*. Allen Lane.

Millikan, Ruth Garrett (1984) *Language, Thought and Other Biological Categories*. MIT Press.

Millikan, Ruth Garrett (1989) 'Biosemantics', *Journal of Philosophy* **86**: 281–97. doi:10.2307/2027123.

O'Connor, Cailin (2019) *The Origins of Unfairness: Social Categories and Cultural Evolution*. Oxford University Press.

Richerson, Peter J and Robert Boyd (1998) 'The Evolution of Human Ultrasociality', in Irenäus Eibl-Eibisfeldt and Frank Kemp Salter, eds., *Ideology, Warfare, and Indoctrinability*: 71–95. Berghahn Books.

Richerson, Peter J and Robert Boyd (1999) 'Complex Societies: The Evolutionary Origins of a Crude Superorganism', *Human Nature* **10**: 253–89. doi:10.1007/s12110-999-1004-y.

Richerson, Peter J and Robert Boyd (2001) 'The Evolution of Subjective Commitment to Groups: A Tribal Instincts Hypothesis', in Randolph M Nesse, ed., *Evolution and the Capacity for Commitment*: 186–220. Russell Sage Press.

Richerson, Peter J, Robert Boyd, and Joseph Henrich (2003) 'Cultural Evolution of Human Cooperation', in Peter Hammerstein, ed., *Genetic and Cultural Evolution of Cooperation*: 373–404. MIT Press.

Ruse, Michael (1986) *Taking Darwin Seriously*. Blackwell.

Seabright, Paul (2010) *The Company of Strangers: A Natural History of Economic Life*, 2nd edition. Princeton University Press.

Shea, Nicholas' (2018) *Representation in Cognitive Science*. Oxford University Press.

Sosis, Richard (2006) 'Religious Behaviors, Badges, and Bans: Signaling Theory and the Evolution of Religion', in Patrick MacNamara, ed., *Where God and Science Meet*, vol. **1**: *Evolution, Genes and the Religious Brain*: 61–86. Praeger.

Sosis, Richard and Candace Alcorta (2003) 'Signaling, Solidarity, and the Sacred: The Evolution of Religious Behavior', *Evolutionary Anthropology* **12**: 264–74. doi:10.1002/evan.10120.

Sperber, Dan (2001) 'An Evolutionary Perspective on Testimony and Argumentation', *Philosophical Topics* **29**: 401–13. doi:10.5840/philtopics2001291/215.

Sperber, Dan, Fabrice Clément, Christophe Heintz, Olivier Msaccaro, Hugo Mercier, *et al.* (2010) 'Epistemic Vigilance', *Mind and Language* **25**: 359–93. doi:10.1111/j.1468-0017.2010.01394.x.

Stanford, P Kyle (2018) 'The Difference Between Ice Cream and Nazis: Moral Externalization and the Evolution of Human Cooperation', *Behavioral and Brain Sciences* **41**: e95. doi:10.1017/S0140525X17001911.

Sterelny, Kim (2012) *The Evolved Apprentice*. MIT Press.

Sterelny, Kim (2021a) *The Pleistocene Social Contract: Culture and Cooperation in the Hominin Lineage*. Oxford University Press.

Sterelny, Kim (2021b) 'Ethnography, Archaeology, and the Late Pleistocene', *Philosophy of Science* **89**: 415–33. doi:10.1017/psa.2021.42.

Sterelny, Kim and Ben Fraser (2017) 'Evolution and Moral Realism', *British Journal for the Philosophy of Science* **68**: 981–1006. doi:10.1093/bjps/axv060.

Stibbard-Hawkes, Duncan N E (2025) 'Reconsidering the Link Between Past Material Culture and Cognition in Light of Contemporary Hunter-Gatherer Material Use', *Behavioral and Brain Sciences* **48**: e1. doi:10.1017/s0140525x24000062.

Timms, Ryan and David Spurrett (2023) 'Hostile Scaffolding', *Philosophical Papers* **52**: 1–30. doi:10.1080/05568641.2023.2231652.

Tomasello, Michael (2016) *A Natural History of Human Morality*. Harvard University Press.

Vavova, Katia (2015) 'Evolutionary Debunking of Moral Realism', *Philosophy Compass* **10**: 104–16. doi:10.1111/phc3.12194.

Whitehouse, Harvey (2008) 'Cognitive Evolution and Religion: Cognition and Religious Evolution', *Issues in Ethnology and Anthropology* **3**: 35–47. doi:10.21301/eap.v3i3.2.

Whitehouse, Harvey (2021) *The Ritual Animal: Imitation and Cohesion in the Evolution of Social Complexity*. Oxford University Press.

Whitehouse, Harvey and Jonathan A Lanman (2014) 'The Ties That Bind Us: Ritual, Fusion, and Identification', *Current Anthropology* **55**: 674–95. doi:10.1086/678698.

Williams, George C (1966) *Adaptation and Natural Selection*. Princeton University Press.

Witteveen, Joeri (2021) 'Biological Markets, Cooperation, and the Evolution of Morality', *The British Journal for the Philosophy of Science* **72**: 401–30. doi:10.1093/bjps/axz007.

GAVIN DAVID YOUNG Lectures *in Philosophy*

ISSN 2653-5378

Peer reviewers Kevin Brosnan (VCU)
Justin Clark-Doane (Columbia University)
Ron Mallon (Washington University in St. Louis)
Editor Antony Eagle (Adelaide)