



Toolmaking and the evolution of normative cognition

Jonathan Birch¹

Received: 5 June 2020 / Accepted: 17 December 2020
© The Author(s) 2021

Abstract

We are all guided by thousands of norms, but how did our capacity for normative cognition evolve? I propose there is a deep but neglected link between normative cognition and practical skill. In modern humans, complex motor skills and craft skills, such as toolmaking, are guided by internally represented norms of correct performance. Moreover, it is plausible that core components of human normative cognition evolved as a solution to the distinctive problems of transmitting complex motor skills and craft skills, especially skills related to toolmaking, through social learning. If this is correct, the expansion of the normative domain beyond technique to encompass more abstract norms of fairness, reciprocity, ritual and kinship involved the elaboration of a basic platform for the guidance of skilled action by technical norms. This article motivates and defends this “skill hypothesis” for the origin of normative cognition and sets out various ways in which it could be empirically tested.

Keywords Normative cognition · Skill · Cognitive control · Norms · Evolution

The skill hypothesis

We are all guided by thousands of norms, often without being able to articulate the norms in question. A “norm”, as I will use the term here, is just any standard of correct or appropriate behaviour. For example, when I take a seat in a lecture theatre, I tacitly choose from a range of appropriate options. I know, without needing to be told, that it would be inappropriate for everyone to sit in the back row, or for everyone to sit directly behind or in front of someone else when there are empty seats with unobstructed views. Yet I could never verbally express the full set of norms that guides my behaviour.

✉ Jonathan Birch
j.birch2@lse.ac.uk
<http://personal.lse.ac.uk/birchj1>

¹ Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK

In such cases, three key psychological ingredients are in place (Railton 2006; Sripada and Stich 2006). First, the agent reliably *notices* or *anticipates* failures to comply with the norm, in themselves or in others. Second, the agent feels *affective pressure* (for example, in the form of discomfort, shame or anger) to prevent or correct the departure from the norm. Third, the agent *knows what to do* to restore conformity in a way appropriate to the situation. This may involve correcting their own behaviour, correcting another's behaviour, asking for forgiveness, or administering punishment. These ingredients are the core elements of *normative cognition*. In short, normative cognition is the micro-regulation of one's own behaviour and the behaviour of others to maintain conformity with norms.

Normative cognition presents a big evolutionary question: *Why did we come to be like this?* Why do norms guide everything we do? How did norms come to structure all human interaction? There is at least a large difference of degree here between humans and all other species, and perhaps a qualitative psychological difference (Sripada and Stich 2006; Machery and Mallon 2010; Kitcher 2011).

I aim here to make progress on this big question by presenting a package of ideas I call the *skill hypothesis*. The first part of the skill hypothesis proposes a psychological connection between normative cognition and certain types of practical skill:

In modern humans, complex motor skills and craft skills, such as toolmaking, are guided by internally represented norms of correct performance.

The second part of the skill hypothesis proposes that this psychological connection indicates a deep evolutionary connection between normative cognition and practical skill:

The capacity to internally represent action-guiding norms of correct performance evolved as a solution to the distinctive problems of standardizing, learning and teaching complex motor skills and craft skills, especially skills related to toolmaking.

In the next three sections, I motivate the skill hypothesis by showing how well it meshes with recent trends in the psychology of skill (Sections 2, 3), by articulating its distinctive empirical predictions (Section 4), and by showing how it yields an empirically credible account of the origin of normative cognition (Section 5). I then discuss how this account relates to others that connect the origin of normative cognition to mutualistic cooperation, altruism and punishment (Section 6). In the last section, I suggest some ways of testing the evolutionary part of the hypothesis against the archaeological record (Section 7).

The underlying motivation for the skill hypothesis is the idea that the central, unifying theme in human cognitive evolution is the cultural transmission of skill within and across generations (Sterelny 2012a). Yet most work on normative cognition has not explicitly connected it to the evolution of skill. My guiding thought is that we will not understand the basic cognitive capacities involved in normative cognition, or their evolution, until we understand the role they play in regulating skilled action. The skill hypothesis emphasizes the role of normative cognition in

craft skill and complex motor skill and postulates a close link between the origin of normative cognition and skilful toolmaking. Precursors to the skill hypothesis (foreshadowing some but not all its elements) include Churchland (1996, 2000); Clark (2000) and Sterelny (2012a).

The cognitive dimension of complex skills

Skill has often been portrayed, in philosophy, psychology and popular culture, as a “mindless” phenomenon in which performance, especially at the expert level, proceeds automatically without any role for cognitive control (Fitts and Posner 1967; Dreyfus and Dreyfus 1986; Dreyfus 2005; Schmidt and Wrisberg 2008; Di Nucci 2013). The commitments of such a view are not easy to spell out, since there is no agreement on what is meant by “automatic” or “cognitive control” (Fridland 2017). A mild version of the “mindless” view simply posits that, when executing an expert-level skill, there is no time for every split-second decision to be preceded by *explicit conscious reflection and deliberation*. I will call this the “weak mindless” view, and I endorse it. A more radical version of the view posits that there is no room for *cognition of any kind* in the successful execution of expert-level practical skills: skill execution is cognitively impenetrable. I will call this the “strong mindless” view, and I reject it.

Recent trends in the psychology and philosophy of skill strongly suggest that the strong mindless view of skill is false (Christensen et al. 2014, 2015a, b, 2016; Pacherie 2008; Papineau 2013; Pavese 2015, 2016a, b; Rietveld 2008; Sutton 2007, 2013; Stanley and Krakauer 2013; Toner et al. 2015; Montero 2016; Stanley and Williamson 2017). Moreover, it is strikingly at odds with the way experts at complex craft skills describe their own expertise. Consider, for example, the following report by the archaeologist Peter Hiscock, a skilled flint knapper:

Knapping long and/or regular sequences is intrinsically a complex process that requires competency at a number of levels simultaneously: bio-mechanical capacity to strike accurately and forcefully; the capacity to anticipate and identify emerging problems in the specimen morphology and to apply an effective action from a repertoire of potential responses; the capacity to plan ahead, which involves mental projections of both future actions and predicted outcomes (Hiscock 2014, p. 34).

More recent theories of skill allow that expert-level practical skill is not wholly automatic. Although there are still many points of disagreement, there is a picture beginning to emerge from this body of work. It is a picture of at least some skills—in particular, complex motor skills and craft skills such as toolmaking—as possessing a substantial role for cognitive control. I propose that along with this role for cognitive control comes a role for normative cognition.

What, then, is the role of cognitive control in expert-level skill execution? Recent literature has tended to focus on sport and dance rather than on craft skill, and the differences are potentially important. In particular, craft skills can typically be

stopped and started at will, left for a while and picked up again later. The time pressure on the skilled agent is less extreme. It is perhaps more plausible, on the face of it, that cognitive control would play a significant role in craft skill than in sport or dance. But I want to start with an example from sport and then draw connections to the case of toolmaking.

Christensen et al. (2015a) focus on elite mountain biking. Their focal example is Kath Bicknell, an expert mountain biker. In advanced mountain biking competitions, competitors must adapt quickly to the challenges of riding an unfamiliar and difficult course, sometimes on an unfamiliar bike: a challenge that requires rapidly modifying one's technique in the face of anticipated, idiosyncratic situational demands. Opportunities for practice runs are severely limited. Information about the course (for example, information about the location of upcoming obstacles such as rocks, and the best strategies for negotiating them) is obtained from a variety of sources: observation of the course, conversations with other bikers, and a handful of practice runs. All of this information (observational, testimonial, episodic) must be integrated with fluent skill execution. For example, Bicknell, anticipating a small rock around the next corner, having remembered it from a practice run, will very slightly adjust her steering and braking to navigate around it.

It is implausible that this integration of situation-specific information with performance is achieved without the involvement of any higher cognitive processing. It is equally implausible that it is achieved through explicit planning, deliberation and reasoning during skill execution, when decisions must be made in tiny fractions of a second. Something intermediate between these extremes is involved: a mode of action control open to new information from a variety of sources, yet also closely integrated with fine motor execution.

Christensen et al. hypothesize that a skilled agent's control of the bike on an unfamiliar course is achieved by means of a cognitive, model-based representation: a *cognitive control model*. A cognitive control model is a representation of the causal structure of a complex skill and the situation in which it is executed. The model mediates between explicit plans and low-level (cerebellar) motor control, representing "causal relations among performance parameters" that allow the individual "to flexibly and appropriately identify and influence [...] parameters in a particular situation", in such a way as to "achieve key performance goals, such as smooth riding and positioning for upcoming obstacles" (Christensen et al. 2015a, p. 344). To visualize intuitively the content of such a model, imagine a causal graph, linking direct handles of agential control (pressure on a pedal, distribution of the rider's weight) to the expected effects on performance (the speed of the bike, the trajectory taken around a corner) given a particular situation (the gradient, the surface, the upcoming obstacles).

Christensen and colleagues' account falls within in a wider trend of positing model-based representations to explain motor control, going back to Conant and Ashby's (1970) "good regulator theorem". For more than two decades, it has been widely thought that motor control relies on "internal models" located in the cerebellar cortex (Wolpert and Kawato 1998, Kawato 1999, Wolpert and Ghahramani 2000, Frith 2012; Blakemore et al. 2000a, b, 2001; Franklin and Wolpert 2011; Pickering and Clark 2014). For example, when one grasps and lifts a cup, the gripping force

applied by the fingers is precisely controlled to be just enough to hold the object in place at each moment. This precise control can be explained by positing an “inverse model” that computes the motor commands needed to achieve a desired result, a “forward model” that predicts future arm trajectories given motor commands, and a “grip controller” that computes required gripping forces given predicted arm trajectories (Kawato 1999). What Christensen et al. posit is a higher level of model-based control, probably dependent on the cerebral cortex rather than the cerebellum, in which a complex skilful activity embedded in a particular situation is modelled, allowing fine adjustment of technique to the details of the situation.^{1, 2}

Like internal models, cognitive control models are *predictive*: they predict the way an action will unfold, given the current set of situational parameters, allowing comparisons between prediction and actual performance, with mismatches between prediction and execution informing future adjustments. In this sense, the proposal in this section is broadly allied with the “predictive processing” movement (as documented in Clark 2015), but without being committed to any wider generalizations about the role of predictive processing in the brain. However, in contrast to cerebellar internal models, the agent has some degree of cognitive access to the parametric structure of the model and is able to make top-down, on-the-fly adjustments to the causal structure and parameter values.

For example, when faced with an unfamiliar course, a skilled biker exploits information from multiple sources to *calibrate* her cognitive control model to the parameters of *this* bike and *this* course, such as the softness of the suspension or the stopping power of the brakes (Christensen et al. 2015a, pp. 347–348). A well-calibrated cognitive control model then guides skill execution in high-performance runs. For instance, Bicknell, having integrated the upcoming obstacle into her control model of the situation, and having calibrated the model to this particular bike, will make a precise adjustment on the brakes well ahead of the obstacle to slow the bike, allowing a line around the corner that evades the obstacle.

As noted above, mountain biking is not toolmaking. However, toolmaking at the level of Acheulean lithic technology or beyond³ presents an agent with the same fundamental challenge: that of adapting a learned technique in the face of anticipated, emerging problems that are unique to *this* situation. Acheulean technology,

¹ There is a welcome convergence here with recent work in moral psychology, which has highlighted the importance of forms of model-based learning in norm acquisition (Cushman 2013; Crockett 2013; Cushman et al. 2017; Railton 2014, 2017a,b; Greene 2017). We are concerned here with model-based *action control* rather than model-based *learning*, but these ideas are complementary: a model-based account of skill execution helps explain how the evaluative models posited by model-based accounts of norm-learning become integrated with fluent agency.

² The general idea of positing intermediate-level predictive models, cognitively accessible to a greater degree than internal models, as a means of explaining the flexibility of proximate control, is originally due to Pacherie (2008), but Christensen et al. apply the idea in detail to complex motor skills.

³ Acheulean toolmaking corresponds to mode 2 in the “five modes” framework introduced by Clark (1969). Later modes are even more cognitively demanding on the toolmaker. For discussion of mode 1 (Oldowan) techniques, see the section “The standardization of technique”. Oldowan toolmaking is usually regarded as much less demanding, but some have argued that we tend to underestimate its sophistication (e.g. Morgan et al. 2015).



Fig. 1 An Acheulean handaxe. Reference: SUSS-64EE9A. Photograph by the Sussex Archaeological Society, via the Portable Antiquities Scheme. CC-BY 2.0

exemplified by the bifacial stone tools usually known as handaxes and cleavers (although their precise use is disputed), is thought to have been first developed by early African *Homo erectus* and was used extensively by *Homo heidelbergensis*, the species often seen as the likely common ancestor of modern humans and the Neanderthals. Acheulean bifaces, especially late examples, show remarkable craftsmanship and symmetry (Lycett 2008; Iovita et al. 2017). Their manufacture was highly skilled and is usually taken to indicate a capacity for the precise guidance of action by socially acquired knowledge (Shipton 2010; Shipton and Nielsen 2015; Shipton 2019).⁴ We see, in Acheulean bifaces, evidence that hominins could make flexible adjustments to a learned technique in response to anticipated, emerging problems. To craft such an object, one must be attuned to the idiosyncratic demands imposed by the shape, size and structure of *this particular piece of stone*, and one must be able to anticipate the problems it will pose (Fig. 1).

I hypothesize that, just as in skilled mountain biking, this on-the-fly modification of technique is made possible by a cognitive control model: a model that represents the control parameters of the process and the likely downstream consequences of small adjustments to technique, allowing extremely consistent generation of a highly symmetric end-product despite wide variety in the situational demands imposed by the stone. This idea is supported by neuroscientific evidence from Stout et al. (2015) and Putt et al. (2017), who show that Acheulean tool production engages mechanisms of cognitive control. Stout and collaborators have long argued for the

⁴ There is a dissenting view that is sceptical of a substantial role for social learning in explaining Acheulean toolmaking (Davidson 2002; Corbey et al. 2016; Tennie et al. 2017). Defenders of this dissenting view may still accept the core elements of the skill hypothesis, but they will regard the forms of toolmaking that indicate shared technical norms as appearing later in the archaeological record.

importance of the Acheulean in human cognitive evolution (but not normative cognition in particular). Their results fit well within the framework of the skill hypothesis (Stout et al. 2008; Stout 2010, 2011; Stout and Chaminade 2012). Along similar lines, Sterelny (2012b) has argued that Acheulean handaxe manufacture was guided by “behaviour programs” and “inner templates” of the end product. I suggest these programs/templates can be assimilated to the notion of a cognitive control model.

This theory of complex motor skills and craft skills as actions guided by well-calibrated cognitive control models is summarised in the following initial account:

Agents who possess a complex motor skill or craft skill possess a well-calibrated cognitive control model that accurately represents those aspects of the causal structure of the situation relevant to successful execution of the skill, *anticipates* upcoming obstacles and problems, and *predicts* the flow of sensory feedback that will occur if skill execution is successful.

The affective dimension of complex skills

Norms are not yet in the picture. This, however, is only a starting point. The above proposal, which emphasizes the cognitive dimension of skill, omits the *affective* dimension. Skill leads to *discontent* when the agent falls short of the standard of performance implicitly encoded in the control model. An incorrect adjustment, leading to a mismatch between the predictions of the cognitive control model and the agent’s behaviour, *feels wrong* to the agent, independently of (and often temporally prior to) any physical discomfort the error may cause. Skill creates internal pressure to conform to an internalized standard of correct performance.

Rietveld’s (2008) work on “situated normativity” is insightful in this respect. Wittgenstein, in the *Lectures on Aesthetics* (1966), urges us to move away from visualizing the central case of aesthetic judgement as that of a passive observer looking at an artwork and remarking, “That’s beautiful”. He asks us instead to imagine a skilled architect inspecting a building under construction at his direction, looking at a doorframe and exclaiming, “Too low! Make it higher!” Wittgenstein’s architect is moved by a reaction of discontent on seeing the doorframe, and this reactive emotion is coupled with immediate knowledge of what must be done to remove it.

Wittgenstein refers to this phenomenon as *directed discontent*: discontent directed towards an object, motivating a specific action to be performed on that object. Rietveld argues that directed discontent is a characteristic feature of all craft skill, not just that of a skilled artist or architect (see also Buskell 2015). It is unobvious only because it often occurs tacitly: Wittgenstein’s case, in which the architect makes an explicit exclamation, is an exception arising from the fact that he cannot personally implement the changes he knows to be required. In most cases, a skilled craftsperson will see a problem, feel directed discontent (some aspect of the situation will *feel wrong*), and make the necessary adjustments without explicit reflection, deliberation or speech. To possess a craft skill is, in part, to feel directed discontent when action falls short of what is required by the situation, and to feel moved to make the required improvements.

Directed discontent can be incorporated within a model-based approach to the psychology of skill by positing that cognitive control models are *integrated with affect*: they are such that certain types of mismatch between the actual performance and the predictions of the model are apt to trigger an emotional response, typically in the form of discontent, directed at the model parameters responsible for the mismatch. When a certain type of mismatch between prediction and execution would, if made, trigger affective pressure directed at the aspect of performance responsible for the mismatch, we can describe this mismatch as an *error of performance by the agent's lights*.

Now we are closing in on norms. For there will not be just one type of mismatch that, if made, triggers affective pressure to modify one's technique. There will be a whole *pattern* of such mismatches. For an elite biker, approaching a corner too fast will feel wrong, motivating adjustment; too much pressure on the brake will feel wrong, motivating adjustment; steering that is too tight or too loose will feel wrong, motivating adjustment—and so on. Skill execution must take a very specific course (the skill must be executed “just the right way”) to avoid triggering any dissatisfaction. This pattern of mismatches implies a *standard of correct performance by the agent's lights*. A cognitive control model implicitly encodes a standard of correct action. When the agent's behaviour falls short of the implicit standard, the performance feels wrong and affective pressure to make an adjustment is triggered. Moreover, the more expertise a person has, the more exacting their implicit standard will be. In an expert, even tiny deviations from the right way of executing the skill cause the performance to feel wrong.

For example, we can think of a skilled craftsperson, such as a flint knapper, as tacitly guided by the norms of their craft. Even if the knapper does not have explicit, conceptual representations of these norms, has no prior grasp of them, and cannot articulate them, the *norms may be encoded implicitly in the pattern of mismatches between model prediction and actual behaviour that would, if made, make the performance feel wrong to the agent*. These mismatches might be (as in the biking example) at the level of fine motor execution: a particular strike of the stone might feel wrong, as the stone fragments in an unexpected way. But in a complex craft skill like flint knapping, they may also occur at a higher, more zoomed-out level: the individual strikes of the flint may feel right in isolation, but the way the specimen is taking shape across a series of successive strikes may feel wrong.

Taking all this into account, here is a proposal that summarises both the cognitive and affective dimensions of skill execution:

Agents who possess a complex motor skill or craft skill possess a well-calibrated cognitive control model that accurately represents those aspects of the causal structure of the situation relevant to successful execution of the skill; *anticipates* upcoming obstacles and problems; *predicts the flow of sensory feedback* that will occur if skill execution is successful; creates *affective pressure* to respond to mismatches between prediction and performance by adjusting one's technique; and represents a *norm of correct performance* in the pattern of mismatches that trigger affective pressure to make an adjustment.

Testing the role of norms in skill execution

What sort of evidence would tell for or against the hypothesis that expert-level complex motor skills and craft skills are guided by norms in the above sense? It is reasonable here to ask for more than subjective reports from experts, which are inconclusive. We could improve on subjective reports by using structured interviews and questionnaires specifically designed to probe the role played by cognitive control and directed discontent in guiding skill execution. We could even attempt to develop a psychometric scale of “normativity about performance”. But perhaps the most striking evidence for the skill hypothesis would come from experiments.

First, let us consider what sort of experiment would support the model-based view of skill execution presented in Section 2. The model-based view contrasts with the strong mindless view on which expert-level skill execution is wholly automatic and fully independent of cognitive control. However, the lack of agreement as to what is meant by “automaticity” leads to some obscurity as to what it is that the strong mindless view posits and the model-based view denies. I rely here on the idea that one signature of cognitive control, as opposed to cognition-free automaticity, is the appropriate, flexible, one-off adjustment of technique, without extensive practice, in response to an anticipated demand imposed by this specific situation (e.g. what is around the next corner).

When the contrast is set up like this, it is clear that the model-based view and the strong “mindless” view do not diverge when a skill is executed in *easy* conditions (when there are no idiosyncratic situational demands to anticipate) and when the situation is *so* novel that the agent is learning a brand new skill, not executing an existing skill (both views predict a role for norms, rules, and so on in these situations). The key, therefore, is to identify “challenging-but-normal” conditions in which the agent is clearly executing an existing skill, but doing so in a demanding situation (e.g. a competitive mountain biking tournament on a difficult course) (Christensen et al. 2016). The strong mindless view predicts that, even in a challenging-but-normal situation, skill execution will be automatic, and thus performance will not be impaired (and might even be enhanced) by engaging the agent’s general cognitive resources in a distracting task (such as remembering lists of numbers). The model-based view predicts that the agent’s performance in such a situation will be impaired by a distracting cognitive task. The prediction is testable but a systematic test has not yet been done. A related prediction of the model-based view is that brain areas associated with cognitive control will be active in challenging-but-normal conditions.

This would test for a cognitive dimension of skill execution, but not specifically for a role for *normative* cognition. This could be put to the test by looking for a differential performance impairment when the distracting task places demands on normative cognition (for example, the distracting task might involve making evaluations of others, or ruling on the severity of norm violations). The prediction is that, controlling for the general cognitive demandingness of the task (as measured when no practical skill is being executed), the effect on skill execution should be more severe when the distracting task is normative than when it is non-normative.

A second, independent testable prediction concerns the familiar phenomenon we might call “*skill dumbfounding*”. It is well known that skilled agents often struggle to articulate the structure of their skills verbally (this can be seen in post-match interviews after any football match), and also well known that people struggle to articulate the reasons for their intuitive normative judgements (Haidt 2001). But are the two abilities correlated across subjects? For example, do people with unusually good access to the structure of their practical skills also have unusually good access to the sources of their normative judgements? The skill hypothesis leads to a general expectation of such a correlation, on the grounds that both kinds of dumbfounding result from the difficulties of translating a norm encoded in a model of a situation into a linguistic format. If no such correlation were found, we could reconcile this with the skill hypothesis by positing differences between the social and technical domains in the level of reflective access to norms (so it would not deal a fatal blow), but finding such a correlation would be positive evidence in favour of the skill hypothesis.

A third type of prediction concerns affective responses to norm violations and failures of skill execution. The skill hypothesis predicts a link between the way we *respond* to failures of skilled action and the way we respond to other types of norm violation. It predicts correlations across subjects: people who feel stronger affective responses to their own failures of skill execution (e.g. strong feelings of shame), should also feel stronger affective responses to their own social norm violations. Likewise, people who react more strongly to others’ violations of technical norms (e.g. with strong feelings of anger) should also react more strongly to others’ violations of social norms. This could be probed experimentally, but it might also be probed through structured interviews with experts.

The evolution of normative cognition

In Sections 2–4, I have hypothesized a link between normative cognition and model-based cognitive control of complex motor skill and craft skill. I will now turn to the second part of the skill hypothesis. If the link is as deep as I have proposed, we are looking at one evolutionary story, not two. Rather than thinking about the evolution of skill and the evolution of normative cognition as wholly separate topics, we should try to understand the evolution of norm-guided skill.

Since the evolution of skill can be traced in the archaeological record, it is possible to construct an empirically credible account. In this section I present an account, informed by archaeological evidence, of how some core components of normative cognition may have evolved.

The great ape platform

Which ingredients of normative cognition were already present in the last common ancestor of *Pan* and *Homo*? On the assumption that the last common ancestor was chimpanzee-like, studies of chimpanzees suggest that social emotions, such as anger

and proto-shame, played a role in managing dominance hierarchies. Anger signalled a credible threat of retaliation in response to insubordination; proto-shame signalled submission to social superiors or the recognition of diminished rank (Fessler 1999; Maibom 2010). Early hominins⁵ kept track of social rank and conditioned their behaviour on rank.

However, this is not guidance of behaviour by norms in anything like the human sense. There are at least large quantitative differences here, perhaps qualitative differences. I will assume in what follows that the last common ancestor did *not* fluently self-regulate their behaviour to maintain conformity with socially learned, culturally variable norms. A human observer, looking in on these early hominins, might *attribute* primitive social norms, such as “don’t take food from a dominant”, but I assume the hominins themselves did not have cognitive control models in which these norms were represented.

There are, in chimpanzees, norm-like phenomena in the form of between-group differences that persist across generations due to social learning. These differences are sometimes called traditions. For example, chimpanzees in some populations display a distinctive “handclasp” behaviour when grooming each other. The style of the handclasp varies between groups, and chimpanzees learn and conform to the prevailing style of their group (van Leeuwen et al. 2012). In one group, a high-ranking female started wearing a piece of grass on her ear, and other group members copied her (van Leeuwen et al. 2014). There are also between-group differences in tool preferences. Some chimpanzee groups prefer to crack nuts with stone hammers, others with wooden hammers, and groups also vary in the size of wooden hammer they prefer (Luncz et al. 2012).

These cultural traditions are sometimes described as norms (e.g. Andrews 2009). However, there is no strong case for positing normative cognition to explain between-group differences. If the high-ranking individuals in a group prefer (for example) stone hammers, and juveniles differentially attend to high-ranking individuals, they may well come to associate the use of stone hammers with reward, and so acquire the same preference by model-free associative learning (Schlingoff and Moore 2017). Stable between-group differences alone do not imply normative cognition.

Chimpanzees’ reactions to infanticide are also sometimes suggested as evidence for norms (Rudolf von Rohr et al. 2011, 2015). Chimpanzees find these events extremely salient, and their attention is drawn strongly to videos of such events featuring unfamiliar individuals, in comparison to videos of hunting or aggression (Rudolf von Rohr et al. 2015). For Rudolf von Rohr et al. (2015), this attentional effect is some evidence of a norm against infanticide. But here too, there is no strong case for positing normative cognition. As the authors note, these are rare events that provoke loud screaming from infants and waa barks from parents, and they would be extremely salient events whether or not the bystanders had internalized a norm against infanticide. Although the comparison videos (of hunting and aggression)

⁵ The term “hominin” encompasses the *Homo* genus and the australopithecines; in other words, modern humans and all ancestral species following the split from *Pan*.

also contained salient features, it was nonetheless the case that “both waa barks and infant screaming were only present in the Infanticide condition” (Rudolf von Rohr et al. 2015, p. 152).⁶

The search for norm-like phenomena in chimpanzees is on-going, and I am not suggesting we can conclude that normative cognition is definitely absent. If it does turn out to be present, the great ape platform I assuming here will turn out to be more minimal than necessary, and some of the steps in the evolutionary account developed here will need to be pushed back in time. Yet even in that scenario, I still think it will have been a useful exercise to construct and test hypotheses about the evolution of hominin normative cognition that begin from a very basic platform.

In addition to social emotions, rank sensitivity, simple cultural traditions, and attentional biases, there were basic capacities for skill learning (including social learning of skills) and high-precision skill execution involving fine motor control (see, e.g., Boesch 2013, Ch. 6, on nut-cracking). A capacity for motor action guided by internal models, largely encapsulated from cognition, may be a core function of the cerebellum with a deep evolutionary history, common to the primates and perhaps to all mammals (Wolpert and Kawato 1998). When a monkey swings between tree branches, it is plausibly guided by an internal model. But did the last common ancestor also possess cognitive control models, allowing top-down modifications of technique in response to anticipated and emerging problems? What is the behavioural marker of this more sophisticated form of action control?

It helps here to return to the example of elite mountain biking: here we find agents making on-the-fly adjustments to a learned technique in response to anticipated challenges that are *idiosyncratic to this situation* (they are challenges posed by *this* bike on *this* terrain, and what is coming around *this* corner). This is also a core feature of craft skill. It is this capacity for flexible adjustment to anticipated challenges unique to the present situation that we took to be evidence of cognitive control models. Do we find this in chimpanzees?

Arguably the most sophisticated learned technique mastered by chimpanzees is ant dipping, in which chimpanzees use tools (such as short poles fashioned from the shrub *Alchornea hirtella*) to harvest army ants. Primatologists have documented several different methods of ant dipping, and chimpanzees are able to adjust their technique to the demands of the situation in a coarse-grained way (reviewed in Humle 2010). Key aspects of technique, such as the length of the tool, are adjusted to situational variables such as the species of the prey ant, and other aspects, such as the method of eating, are in turn adjusted to the length of the tool.

However, what seems to be missing is evidence of flexibility in the face of *anticipated*, as opposed to observed, situational demands. Sensitivity to regular, recurring features of situations, such as the species of prey ant, may be achieved without a

⁶ Rudolf von Rohr et al. (2015) argue that we can exclude the alternative explanation based on salience, but they do not make a compelling case. They note that the difference in looking times between the video with the most infant screams and the fewest was not statistically significant (2015, Table 2), but the absence of a statistically significant difference between these two videos, in a study of 17 animals, is not compelling evidence of the absence of an effect.

cognitive control model by having several discrete techniques available to the motor system, together with the use of perceptual cues to identify the *general type* of situation one is in. The paradigm here is grasping a cup: different types of handle require different grasping techniques, but this coarse-grained sensitivity to general types of situation is not the same as the ability to anticipate emerging challenges and problems. While the evidence is inconclusive, I hypothesize that cognitive control models are unique to the hominin lineage. The caveat noted above applies here too: if cognitive control models turn out to be present in chimpanzees, then the great ape platform I am assuming here will turn out to be more minimal than necessary.

The standardization of technique

Human cognitive evolution is a story of the dramatic expansion and enrichment of skill learning and skill execution capabilities, driven by feedback loops connecting social foraging, social learning and environmental change. Kim Sterelny (2012a) provides an account of how these feedback loops worked. In brief, hunting and gathering in a changing environment put an evolutionary premium on the social transmission of skill. Foragers needed to know what to hunt and gather and how to hunt and gather it, and in a changing environment the necessary skills changed too quickly to be genetically encoded. Under this selection pressure, hominins developed a form of apprentice learning. Early hominins acquired skills from the preceding generation through a process of scaffolded trial-and-error: juveniles would stay close to adults, observe both the performance and the products of their skills (such as completed tools), and attempt to recreate those performances and products. The learning environment was “scaffolded” in the sense that learners were immersed in a community of agents performing and completing the skills they needed to learn, allowing them to see their own errors and bring their internal standards for correct performance into line with those of the group.

These scaffolded learning environments co-evolved with cognitive and life-history adaptations: adults evolved to be tolerant and supportive of apprentices, who evolved to learn from their models during a protracted period of juvenile cognitive development. Each incremental increase in the bandwidth and volume of social learning led to incrementally more efficient social foraging, which in turn allowed for incrementally more protracted learning. Meanwhile, more efficient social foraging further accelerated the pace of environmental change, putting an even greater premium on the social learning of skill.

The long-term results can be seen in the symmetry and precision of Acheulean bifaces. I have already made a case, in Section 2, for thinking of Acheulean toolmaking as guided by a form of model-based cognitive control. This suggests that *H. heidelbergensis* had evolved this form of cognitive control, and could make flexible adjustments to a learned technique in response to anticipated, emerging problems.

This is evidence, then, that the *cognitive* dimension of skill as seen in modern humans was in place by the Acheulean period. However, to be clear, I am not ruling out an earlier date of origin for cognitive control models. Oldowan tool production is much less complex, and provides less persuasive evidence of cognitive control

(Ambrose 2001; Wynn and McGrew 1989), but Morgan et al. (2015) argue that the received wisdom about the Oldowan may have underestimated its complexity.

What about evidence of the *affective* dimension of skill? I have argued that the affective dimension of skill consists in affective pressure that motivates an agent to correct or forestall performance errors by modifying technique. In short, the performance *feels* wrong (or feels right), even before it has manifestly *gone* wrong (or right). We can identify two plausible selection pressures for this affective dimension of skill.

First, note that the manufacture of Acheulean bifaces was, in at least some cases, a collaborative activity involving a division of labour. Shipton and Nielsen (2015) report evidence of spatial division of labour at a site in India, with the early stages and finishing stages of cleaver manufacture carried out at distinct locations. The best explanation, they argue, is that different group members were undertaking different tasks. Inherent to collaborative tool manufacture is a special kind of coordination problem resulting from the causal opacity of complex skills: individual agents at the early stages of the process will not be fully aware of what the finishers do, and consequently will not be fully aware of the downstream consequences of their actions. Small variations of technique in the early stages may result in the finisher receiving a tool that cannot be finished. The consequences of these variations will not be readily foreseeable unless the agents at the beginning of the production line have also mastered the skills to be performed later, which would undercut the advantages of dividing the labour.

A solution to this problem is the *standardization of technique*: all agents involved in the process internalize a particular way of doing their part and conform to that way even when departures from it seem, from their point of view, inconsequential. There is archaeological evidence that techniques were indeed standardized at particular sites (Wynn 1993; Shipton 2010). A recent study of late Acheulean handaxe types in Britain found a range of distinct subtypes at different sites, leading the authors to propose that that “the distinctive and difficult to produce handaxes types that characterize the British Late Acheulean were reproduced according to normative expectations of what handaxes should look like” (Shipton and White 2020, p. 1). These standardized techniques were passed down the generations with high fidelity.

In explaining how they were initially passed down, we should not underestimate the importance of the physical tools themselves, in either complete or unfinished form, which provide a physical template (an “outer template”, as it were), for what constitutes successful toolmaking. Each generation would have internalized the standards of the preceding generation by observing their technique and by learning to copy those tools to which they assigned particular value. I see no opposition between positing cognitive control models internal to the agents and positing that physical tools also provided valuable *external* models. Both can be true: the physical tool, finished or unfinished, can provide a benchmark against which the agent’s internal model of correct skill execution is calibrated.

Standardization of a complex skill indicates not just cognitive control, but also a robust motivation to adhere to one particular style of skill execution among multiple possible styles. It involves *internalizing* a particular style of skill execution, in the

sense of feeling negative affective responses to small departures from that style. I suggest that standardized toolmaking is an archaeological marker of norm-guided skill execution. Agents were, I suggest, internalizing technical norms—*our ways* of executing practical skills—and feeling negative affect when their performance failed to meet the group's shared standards.⁷

Moreover, I hypothesize that standardization is not just a marker but also a *driver*: the benefits of group-wide standardization of toolmaking technique created a selection pressure that favoured (groups of) agents who felt negative affective responses when they deviated from a cognitive control model. Although this is a selection pressure that arises in the context of a group, the benefit here is a direct fitness benefit: agents who internalize group-wide standards themselves benefit, over the long run, from more efficient toolmaking.⁸

What is the second selection pressure? Mastering a skill as complex as Acheulean tool manufacture requires years of sustained practice. This marks a significant watershed compared to anything observed earlier in the archaeological record. What motivates sustained practice? Mastering these skills would have yielded long-term direct fitness benefits, especially if high quality tools could be exchanged for other resources, but it is implausible to suppose that explicit knowledge of the long-term fitness benefits motivated practice. Given the long-term fitness benefits and the difficulty of being motivated by them, I hypothesize that selection favoured agents *intrinsically motivated to master skills*—agents who felt satisfaction in achieving excellence in skill execution, and discontent at any aspect of performance that fell short of their internalized standard.

Monitoring and teaching others

Standardized techniques of tool manufacture were “norms” in a broad sense: they provided a standard of correct performance to which a skilled toolmaker was intrinsically motivated to conform. They were also social in the sense of being coordinated across groups. But they need not have involved third-party monitoring, expectation and/or enforcement, which are common features of modern social norms (Bicchieri 2005). These elaborations were, I suggest, connected to teaching.

It is not clear when intentional, deliberate teaching first emerged. Scaffolded trial-and-error learning, of the type described by Sterelny (2012a, b, c), can occur without intentional teaching. Models need to tolerate the presence of learners, but they do not have to knowingly demonstrate their skills or instruct their apprentices. Even now, the ethnographic record contains many examples of children learning skills

⁷ In a forthcoming paper, Shipton et al. ([in press](#)) also suggest that standardized toolmaking techniques are an archaeological signature of norms. However, they do not argue that the demands of toolmaking were themselves one of the drivers of the evolution of normative cognition.

⁸ Processes in which change is driven by a direct fitness benefit, but in which the benefit depends on organisms living in groups, are sometimes described as cases of “group selection”, but this is in some ways a misleading term. See Birch (2017, Ch. 4) for discussion of how to label these cases.

without any intentional teaching, by means of free exploration in an environment full of people displaying the skill (Lancy 2016).

However, Hiscock (2014) argues that, given the high risk of injury inherent in even very small errors, Acheulean tool manufacture was probably transmitted through intentional teaching. Shipton (2019, pp. 343-347) argues that this is especially plausible of the Late Acheulean (from 600,000 years ago), due to the presence in the archaeological record of subtle forms of platform preparation that would be extremely difficult to learn without teaching. Teaching, as opposed to mere toleration of learners, relies on comparing the behaviour of *others* to one's own control model: one must notice and anticipate failures of skilled performance in the learner, and one must respond appropriately to these failures by signalling the nature of the failure, its seriousness, and the appropriate corrective measure.⁹

In cases where errors by a learner are foreseen and managed by a teacher, physical injury can be avoided: milder punishments and rewards administered by the teacher can take their place. We see in the evolution of teaching the first stage in the socialization of cognitive control: a good teacher models the actions *of the learner*, identifying mismatches between the learner's performance and the predictions of the model, and responding with corrective measures to errors of performance. Having initially evolved for the self-regulation of behaviour, I hypothesize that control models were co-opted for the regulation of the behaviour of other agents.

A critic might argue that this suggestion faces a chicken-and-egg problem. How plausible is it that pedagogical contexts, in which a learner attends carefully to a model who tolerates their presence and supports their learning, arose *prior* to widespread enforcement of social norms? In contemporary humans, pedagogy is a richly norm-governed activity. We instil in children norms of respect for teachers, and we have a normative expectation that adults will tolerate, and help to correct, the mistakes of children. If norms arose first in the context of toolmaking, and only later expanded outwards to other domains, we have to assume that pedagogical contexts initially arose without the enforcement of pedagogical norms of respect for teachers and tolerance of learners. This is a potential problem for any account which begins with self-regulation in the context of skill execution, moves from there to other-regulation in the context of skill transmission, and from there to more general, more abstract social norms.

However, methods of proto-teaching in the Acheulean would not have placed anything like as much strain on the agents' normative capacities as contemporary methods of teaching. Teaching in modern human societies typically involves groups of learners, taught by teachers who are neither their relatives nor their co-residents. The group size strains the learners' attentional capacities, and cooperating with non-relatives and non-co-residents strains their emotional capacities, leading in turn to strain on the teacher's capacity for tolerance. Norm enforcement is

⁹ Boesch (2013, pp. 143-144) offers two anecdotes of active error correction in Tai chimpanzees. So, even at this stage of the account, it's hard to be sure that the traits being described are uniquely hominin. If they are not, the central idea of an expansion of the normative domain from technical to non-technical norms may still be correct, but we would have to push back the date for the origin of technical norms.

necessary to stabilize the situation. The challenges, for both teachers and learners, of regulating their emotions in the classroom have been well documented by education researchers (Sutton 2004; Graziano et al. 2007; Gullone et al. 2010). By contrast, the transmission of skill in the Acheulean would have occurred between models and learners who were, if not closely related, then at least co-resident in the same camp. Even now, in contemporary hunter-gatherer societies, craft skills are often transmitted primarily from parents to offspring, usually from the same-sex parent (Hewlett and Cavalli-Sforza 1986; Shennan and Steele 1999; Shennan 2002; Mameli 2008). Prosocial dispositions towards campmates and kin would have been enough to allow simple forms of teaching, based on the demonstration of skills, to be accomplished without excessive emotional strain prior to the enforcement of pedagogical norms.

The repurposing of shame

While these gradual transformations were in progress, there was also a transformation in the affective salience of norm violations. In the context of skilled toolmaking, directed discontent plays an important role (a skilled knapper sees errors and is motivated to correct them) but this affective response need not take the form of shame or outrage. As noted in Section 2, proto-shame and anger plausibly originated in dominance hierarchies, where they regulate social rank (Fessler 1999); in hominins, these emotions were co-opted for the regulation of prestige hierarchies (Fessler 2004, 2007).

As obligate collaborators, early humans constantly evaluated each other as prospective partners in collaborative enterprises, including toolmaking, and reputation management was crucial for success (Boehm 2012; Baumard 2016; Tomasello 2016). Initially, these evaluations need not have involved normative standards: they may simply have been predictive judgements about how likely a potential partner is to succeed at some task. Agents would prefer cooperative partners with a track record of success. But as normative cognition became more sophisticated for the reasons suggested above, these evaluations would have acquired a normative dimension. Once agents were monitoring each other's behaviour for conformity with an internalized standard, it became possible to judge another agent to be doing things *the wrong way*, independently of any manifest, visible failures that resulted from the error.

I hypothesize that shame and anger acquired new functions in this context: they became a means of signalling recognition that one's own performance, or the performance of another, fell *severely* short of the group's shared internalized standards, and a means of motivating improvement (see Tangney et al. 2013 and Sznycer et al. 2016 for related ideas). Falling short of the standard in the presence of others triggers shame in oneself and anger in others, and these affective responses provide a strong motivation to do better next time.

The expansion of the normative domain: fairness, reciprocity, ritual and kinship

At this point in the account, normative cognition remains restricted to a core set of socially learned skills: it is far from a ubiquitous feature of human social life. I hypothesize that this ubiquity was a later development. *The normative domain expanded from technical norms to incorporate norms of fairness, reciprocity, ritual and kinship.* I suggest that these norms, despite their apparently abstract and general character, were an elaboration of a basic capacity for norm-guided skill execution.

How did the process of “elaboration” occur? Simple norms of fairness may have arisen in the context of collaborative hunting: skilful execution of a hunt would flow into skilful division of the spoils, itself guided by technical norms specifying our way of dividing a carcass. Norms of equitable division would have been favoured because they benefited the agent, in the long run, by showing them to be a trustworthy and profitable cooperation partner (Baumard 2016).

Norms of reciprocal exchange may have originated with the emergence of large-scale exchange networks in the late Palaeolithic (Marwick 2003). Sterelny (2014) has argued that expanding social groups and exchange networks created a “Palaeolithic reciprocation crisis”, a package of coordination problems resulting from the demands of reciprocity in large networks. Larger groups favour greater specialization: in a late Pleistocene “tribal” network of 500 or more individuals there might, for example, be market for a full-time specialist toolmaker or spear-thrower. But specialization requires reciprocal exchange (e.g. of tools for food), and reciprocal exchange requires norms of market value: one must know how much food a handaxe is worth, for example. These norms, although apparently quite abstract, may have begun as norms of skilled behaviour in specific situations: norms of how to barter skilfully round the campfire, norms of what to offer and what to accept in one-on-one interaction.

Larger groups also faced the problem of creating group cohesion by means other than one-on-one bonding (Dunbar 2014). A solution was ritual: skilful collective performances, high in emotional and mnemonic resonance. We can see how a capacity to internalize technical norms, including norms concerning long sequences of actions, would bring with it a capacity for ritual. Indeed, toolmaking practices can themselves resemble rituals. Norms of ritual may have begun as norms of skill execution: norms of how to dance, or make music, our way.

Norms of ritual are, in turn, a step towards norms of kinship. At some rituals, inter-band monogamous pairings would have been initiated, guided by norms of who can pair up with whom (Chapais 2008; Allen et al. 2011). On the face of it, kinship norms are among the most abstract norms, specifying (for example) that one may marry a cross-cousin but not a parallel cousin. However, the first kinship norms may have been concrete norms of skill execution in group rituals: a skilled performer knows where to go, whom to dance with, which moiety (i.e. descent group) to attach to, and is led by norms of ritual behaviour to an appropriate mating partner.

In short, behaviourally modern humans evolved elaborate systems of abstract norms encompassing trade, ritual and family life, but I hypothesize that these norms were learned, stored and executed using mechanisms that had originally evolved for the standardization and teaching of toolmaking techniques.

Comparison to other hypotheses

Normative cognition is a complex adaptation, assembled incrementally over many thousands of generations, and any adequate explanation of its evolution must appeal to a sustained selection pressure over a long period of time. The skill hypothesis proposes that the demands of toolmaking provided the pressure. Acheulean and post-Acheulean toolmaking is a complex skill with multiple possible styles of execution, calling for (1) sophisticated mechanisms of model-based cognitive control, (2) group-wide standardization of technique through technical norms, (3) years of sustained practice, driven by an intrinsic motivation to achieve an internalized standard of correct performance, and (4) at least a basic form of teaching, in which a teacher monitors the performance of a learner and notices errors. The demands of standardization, practice and teaching selected for cognitive changes that made them easier, enabling yet higher levels of standardization, yet more demanding practice regimes, and more sophisticated forms of teaching. In short, normative cognition was the adaptation that facilitated the standardization, practice and teaching of our ways of performing complex practical skills. This can be regarded as an example of a “gene-culture co-evolution” hypothesis (Boyd and Richerson 1985). That said, I have avoided explicitly assuming that mechanisms of model-based control are genetically inherited. If they turn out to be themselves culturally inherited, then the above story may be reinterpreted as one of “culture–culture co-evolution” (Birch and Heyes 2020).

It is worth comparing this idea with other possible gene-culture co-evolution hypotheses: those which posit a role for forms of mutualistic cooperation other than toolmaking, and those which posit a crucial role for altruism and punishment. There is no big disagreement here about the importance of cooperation for the evolution of normative cognition, or the importance of gene-culture co-evolution (or perhaps culture–culture co-evolution), but there is disagreement about the *type* of cooperation that mattered most.

Other cooperative activities

The skill hypothesis proposes that normative cognition evolved in response to the demands of complex skills that required years of sustained practice and groupwide standardization of technique. Acheulean toolmaking is, I suggested above, the earliest example of a skill with the requisite features. But I remain open to the possibility that other collaborative activities were at least as important, and also possessed the required features. Big game hunting—another demanding collaborative activity requiring at least some division of labour—is another example of a complex motor skill that may have the required features.

However, the skill hypothesis predicts that not *any* form of mutualistic cooperation selects for normative cognition, because not any form of mutualistic cooperation will involve a complex skill with the right features. The skill hypothesis proposes that it was the need to standardize technique and to motivate practice that initially

selected for the internalization of normative standards of correct performance. On this hypothesis, normative cognition evolved, as Tomasello (2016, 2020) and Bauvard (2010) have also suggested, in the context of skilful mutualistic cooperation. But it was not the mutualistic character of the behaviour that generated the selection pressure for technical norms: it was its skilful character. Simpler mutualistic activities would not have created the same selection pressures.

Altruism

Existing approaches often link the evolution of normative cognition to the evolution of altruism (e.g. Joyce 2006; Kitcher 2011). The general idea is that normative cognition evolved to make us more altruistic or to “remedy altruism failures” (Kitcher 2011, p. 135). The skill hypothesis is compatible with the idea that, once early humans had a fairly sophisticated capacity for normative cognition, *one* of the things they could do with it was internalize norms of altruism. What it rejects is the idea that a need to be more altruistic was an initial driver of the evolution of normative cognition.

One problem with the “altruism failures” hypothesis is that many norms, including highly robust norms such as incest prohibitions, have nothing to do with altruism (Sripada 2005; Sripada and Stich 2006; Machery and Mallon 2010). However, many norms have nothing to do with skill either. It is part of the skill hypothesis that the normative domain gradually expanded beyond technical norms to encompass more abstract domains via norms of ritual and collaborative activity. A defender of the altruism failures hypothesis could argue, in a similar vein, that the normative domain gradually expanded beyond altruism failures to encompass non-altruistic norms.

A different problem is that altruism, in the psychologist’s sense of action motivated by concern for others, does not require normative cognition: what it requires is sympathy, bonding and trust (Sterelny 2012c, pp. 103–104). Altruism-based accounts of the evolution of normative cognition thus face the question: *If the problem was a need to boost altruism, why was it not solved by tuning up the intensity of sympathy, bonding and trust?* Selection for variants that dialled up these emotions would have solved the problem of altruism failures more directly and effectively than normative cognition. This too is not a decisive refutation of the “altruism failures” hypothesis. Sometimes natural selection assembles overcomplicated solutions to design problems. But it is unsatisfying to regard normative cognition as nothing more than an overcomplicated way of solving a problem that could have been solved by fine-tuning an existing affective mechanism. The skill hypothesis shows there were other design problems—the need to standardize technique across groups, to motivate sustained practice, and to transmit the shared standards down the generations—to which the distinctive features of normative cognition provide a solution.

Punishment

Some existing approaches also link normative cognition to the evolution of punishment (Sripada 2005; Machery and Mallon 2010; Kitcher 2011, pp. 90–91). The idea

is that normative cognition evolved to help individuals evade punishment from others. As with altruism, the skill hypothesis is compatible with the suggestion that, once early humans had a sophisticated capacity for normative cognition, *one* of the things they could do with it was internalize norms of what to punish and how to punish it. It denies, however, that pre-existing punishment practices were an initial driver of the evolution of normative cognition.

If we are talking about cognitively sophisticated forms of punishment, in which the perpetrator is judged to have violated a norm, then these capacities presuppose normative cognition and cannot pre-date it (Sterelny 2012c, p. 105). If, however, we are talking about very simple punishment-like phenomena, such as retaliation in a dominance hierarchy, we again face the problem that normative cognition would have been an overcomplicated solution to the problem. Model-free associative learning allows agents to learn associations between certain types of action and retaliation. If the problem is insufficient fear of alphas, a direct and effective solution is to dial up fear of alphas. Yes, natural selection sometimes assembles overcomplicated solutions to design problems. But it is an advantage of the skill hypothesis that it does not lead to the same kind of mismatch between the features of normative cognition and the design problem it initially solved.

Language: precondition or consequence?

Some accounts of the evolution of normative or moral cognition assume that language is already on the scene and take their explanatory target to be language-dependent (Gibbard 1990; Joyce 2006). By contrast, language plays no significant role in the skill hypothesis, and the elements of normative cognition the hypothesis aims to explain are language-independent. We do not know when language originated, but there is no reason to suppose the manufacture of Acheulean bifaces required it (Putt et al. 2014). A relationship the other way, whereby skilled toolmaking enabled and drove the evolution of language, is more plausible (Stout and Chaminade 2012; Sterelny 2012b). The same goes for teaching and the group-wide standardization of technique: these things may require signalling, but the signalling can be gestural. The story is compatible with an origin for full language long after the basic elements of normative cognition were in place.¹⁰ This is just as well, since language is itself a richly norm-guided skill, and accounts that assume it to be there at the beginning face a chicken-and-egg problem.

¹⁰ For example, full language may be linked to the emergence of large-scale exchange networks in the late Palaeolithic, described earlier, and the consequent need to keep track of the obligations created by norms of kinship and reciprocity outside of contexts where kin and trading partners were physically present (Marwick 2003). Or perhaps full language evolved to enable gossip, a low-cost mechanism for facilitating norm enforcement and reciprocity (Dunbar 1996). These are just two of the possibilities that come into view once we allow that language post-dates normative cognition.

Looking for Archaeological Evidence

In Section 4, I argued that the first part of the hypothesis—the proposed link between the psychology of norms and the psychology of skills—could be tested empirically. The same applies to the evolutionary part of the skill hypothesis. This part makes claims about the relationship between the evolution of norms and the evolution of skills that can be tested against the archaeological record.

The central empirical prediction is that *technical norms (as shown by standardized toolmaking) should not post-date more abstract norms of ritual, reciprocity, fairness and kinship*. If there are cases where abstract norms unrelated to skill can be shown to pre-date technical norms for specific skills, this would refute the skill hypothesis. It would also count against the hypothesis if there are examples of societies with very simple technical norms but highly complex, highly abstract social norms. Another prediction is that abstract social norms, such as norms of kinship, should be such that they could, at least initially, have been enacted as norms of skilled performance in a specific context, such as a ritual.

A further prediction is that the complexity of the technical and non-technical norms a society can support will be correlated, both in the ethnographic record and over archaeological time. A population's ability to support more complex technical norms will be linked to an ability to support more complex norms in the rest of the social world. A clear step up in the sophistication of technical norms, such as the shift from mode 2 to mode 3 toolmaking, is predicted to be followed in the archaeological record by evidence of a step up in the sophistication of non-technical norms. The question of how closely the expansion of the normative domain tracked the progression through the “five modes” of toolmaking is an important question for further work.

I don't want to exaggerate the support the skill hypothesis would receive from the confirmation of these predictions, since it would no doubt be possible to reconcile the evidence with contrasting hypotheses. My aim is simply to give a sense of what a future archaeological case for the skill hypothesis might look like.

In sum, although it has yet to receive serious empirical attention, I have made a case that the skill hypothesis is supported by recent trends in the psychology of skill and compatible with existing data on the archaeology of skill. By bringing out a neglected connection between the psychology of skill and the psychology of norms, the hypothesis opens up new lines of investigation for cognitive science, archaeology, evolutionary anthropology and philosophy.

Acknowledgements I thank Ali Boyle, Richard Bradley, Liam Kofi Bright, Campbell Brown, Susanne Burri, Andrew Buskell, Wayne Christensen, Ellen Fridland, Mattia Gallotti, Balasz Gyenis, Cecilia Heyes, Peter Hiscock, Laurenz Hudetz, Anton Leist, Peter Railton, Ceri Shipton, David Spurrett, Kim Sterelny, John Sutton, Johanna Thoma, Felix Warneken, Joeri Witteveen and two anonymous reviewers for their comments and advice. I thank audiences at the Australian National University (ANU), the London School of Economics and Political Science, the University of Cambridge, the University of Notre Dame and the University of Utrecht for their excellent questions.

Funding This work was supported by a Philip Leverhulme Prize from the Leverhulme Trust (Grant No. PLP-2014-399) and visiting fellowships at the ANU in 2015 and 2018.

Compliance with ethical standards

Conflicts of interest Author declares that he has no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Allen NJ, Callan H, Dunbar R, James W (eds) (2011) Early human kinship: from sex to social reproduction. Wiley, Hoboken, NJ
- Ambrose SH (2001) Paleolithic technology and human evolution. *Science* 291:1748–1753
- Andrews K (2009) Understanding norms without a theory of mind. *Inquiry* 52:433–448
- Baumard N (2016) The origins of fairness: how evolution explains our moral nature, trans. Reeve P. Oxford University Press, New York
- Bicchieri C (2005) The grammar of society: the nature and dynamics of social norms. Cambridge University Press, Cambridge
- Birch J (2017) The philosophy of social evolution. Oxford University Press, Oxford
- Birch J, Heyes CM (2020) The cultural evolution of cultural evolution. *PsyArXiv*. <https://psyarxiv.com/nmuva/>. Accessed 23 September 2020
- Blakemore S-J, Wolpert D, Frith C (2000a) Why can't you tickle yourself? *NeuroReport* 11:R11–R16
- Blakemore S-J, Wolpert D, Frith C (2000b) The cerebellum is involved in predicting the sensory consequences of action. *NeuroReport* 12:1879–1884
- Boehm C (2012) Moral origins: the evolution of virtue, altruism and shame. Basic Books, New York
- Boesch C (2013) Wild cultures: a comparison between chimpanzee and human cultures. Cambridge University Press, Cambridge
- Boyd R, Richerson PJ (1985) Culture and the evolutionary process. University of Chicago Press, Chicago, IL
- Buskell A (2015) How to be skilful: normative sensitivity and opportunistic robustness. *Synthese* 192:1445–1466
- Chapais B (2008) Primeval kinship: how pair-bonding gave birth to human society. Harvard University Press, Cambridge, MA
- Christensen W, Bicknell K, McIlwain D, Sutton J (2015a) The sense of agency and its role in strategic control for expert mountain bikers. *Psychol Conscious* 2:340–353
- Christensen W, Sutton J, McIlwain DJF (2015b) Putting pressure on theories of choking: towards an expanded perspective on breakdown in skilled performance. *Phenomenol Cogn Sci* 14:253–293
- Christensen W, Sutton J, McIlwain DJF (2016) Cognition in skilled action: meshed control and the varieties of skill experience. *Mind Lang* 31:37–66
- Churchland PM (1996) The neural representation of the social world. In: May L, Friedman M, Clark A (eds) *Mind and morals: essays on cognitive science and ethics*. MIT Press, Cambridge, MA, pp 91–108
- Churchland PM (2000) Rules, know-how and the future of moral cognition. *Can J Philos* 30:291–306
- Clark G (1969) *World prehistory: a new synthesis*. Cambridge University Press, Cambridge
- Clark A (2000) Word and action: reconciling rules and know-how in moral cognition. *Can J Philos* 30:267–289
- Clark A (2015) *Surfing uncertainty: prediction, action, and the embodied mind*. Oxford University Press, New York

- Conant RC, Ashby WR (1970) Every good regulator of a system must be a model of that system. *Int J Syst Sci* 1:89–97
- Corbey R, Jagich A, Vaesen K, Collard M (2016) The Acheulean handaxe: more like a bird's song than a Beatles' tune? *Evol Anthropol* 25:6–19
- Crockett MJ (2013) Models of morality. *Trends Cogn Sci* 17(8):363–366
- Cushman F (2013) Action, outcome, and value a dual-system framework for morality. *Pers Soc Psychol Rev* 17(3):273–292
- Cushman F, Kumar V, Railton P (2017) Moral learning: psychological and philosophical perspectives. *Cognition* 167:1–10
- Davidson I (2002) The 'finished artefact fallacy': Acheulean handaxes and language origins. In: Wray A (ed) *The transition to language*. Oxford University Press, Oxford, pp 180–203
- Di Nucci E (2013) *Mindlessness*. Cambridge Scholars Publishing, Cambridge
- Dreyfus HL (2005) Overcoming the myth of the mental: how philosophers can profit from the phenomenology of everyday expertise. *Proc Address Am Philos Assoc* 79:47–65
- Dreyfus HL, Dreyfus SE (1986) *Mind over machine: the power of human intuition and expertise in the era of the computer*. Free Press, New York
- Dreyfus HL, Dreyfus SE (1991) Towards a phenomenology of ethical expertise. *Hum Stud* 14:229–250
- Dunbar R (1996) *Grooming, gossip and the evolution of language*. Harvard University Press, Cambridge, MA
- Dunbar R (2014) *Human evolution*. Pelican, London
- Fessler DMT (1999) Toward an understanding of the universality of second order emotions. In: Hilton AL (ed) *Beyond nature or nurture: biocultural approaches to the emotions*. Cambridge University Press, Cambridge, pp 75–116
- Fessler DMT (2004) Shame in two cultures: implications for evolutionary approaches. *J Cogn Cult* 4:207–262
- Fessler DMT (2007) From appeasement to conformity: evolutionary and cultural perspectives on shame, competition, and cooperation. In: Tracy JL, Robins RW, Tangney JP (eds) *The self-conscious emotions: theory and research*. Guilford Press, New York, pp 174–193
- Fitts PM, Posner MI (1967) *Human performance*. Brooks/Cole Publishing, Belmont, CA
- Franklin DW, Wolpert DM (2011) Computational mechanisms of sensorimotor control. *Neuron* 72:425–442
- Fridland E (2017) Automatically minded. *Synthese* 194:4337–4363
- Frith C (2012) Explaining delusions of control: the comparator model 20 years on. *Conscious Cogn* 21:52–54
- Gibbard A (1990) *Wise choices, apt feelings: a theory of normative judgment*. Harvard University Press, Cambridge, MA
- Graziano PA, Reavis RD, Keane SP, Calkins SD (2007) The role of emotion regulation and children's academic success. *J Sch Psychol* 45:3–19
- Greene JD (2017) The rat-a-gorical imperative: moral intuition and the limits of affective learning. *Cognition* 167:66–77
- Gullone E, Hughes EK, King NJ, Tonge B (2010) The normative development of emotion regulation strategy use in children and adolescents: a 2-year follow-up study. *J Child Psychol Psychiatry* 51:567–574
- Haidt J (2001) The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychol Rev* 108:814–834
- Hewlett BS, Cavalli-Sforza LL (1986) Cultural transmission among Aka pygmies. *Am Anthropol* 88:922–934
- Hiscock P (2014) Learning in lithic landscapes: a reconsideration of the hominid "toolmaking" niche. *Biological Theory* 9:27–41
- Humle T (2010) How are army ants shedding new light on culture in chimpanzees? In: Lonsdorff EV, Ross SR, Matsuzawa T (eds) *The mind of the chimpanzee: ecological and experimental perspectives*. University of Chicago Press, Chicago, IL, pp 116–127
- Iovita R, Tuvi-Arad I, Moncel M-H, Despriée J, Voinchet P, Bahain J-J (2017) High handaxe symmetry at the beginning of the European Acheulian: the data from la Noira (France) in context. *PLoS ONE* 12(5):e0177063
- Joyce R (2006) *The evolution of morality*. MIT Press, Cambridge, MA
- Kawato M (1999) Internal models for motor control and trajectory planning. *Curr Opin Neurobiol* 9:718–727

- Kitcher P (2011) *The ethical project*. Harvard University Press, Cambridge, MA
- Lancy DF (2016) Playing with knives: the socialization of self-initiated learners. *Child Dev* 87:654–665
- Luncz LV, Mundry R, Boesch C (2012) Evidence for cultural differences between neighboring chimpanzee communities. *Curr Biol* 22:922–926
- Lycett SJ (2008) Acheulean variation and selection: does handaxe symmetry fit neutral expectations? *J Archaeol Sci* 35(9):2640–2648
- Machery E, Mallon R (2010) Evolution of morality. In: Doris J (ed) *The moral psychology handbook*. Oxford University Press, Oxford, pp 3–47
- Maibom HL (2010) The descent of shame. *Philos Phenomenol Res* 80(3):566–594
- Mameli M (2008) Understanding culture: a commentary on Richerson and Boyd's *Not By Genes Alone*. *Biol Philos* 23:269–281
- Marwick B (2003) Pleistocene exchange networks as evidence for the evolution of language. *Cambridge Archaeol J* 13:67–81
- Montero BG (2016) *Thought in action: expertise and the conscious mind*. Oxford University Press, Oxford
- Morgan TJH, Uomini NT, Rendell LE, Chouinard-Thuly L, Street SE, Lewis HM, Cross CP et al (2015) Experimental evidence for the co-evolution of hominin tool-making teaching and language. *Nat Commun* 6:1–8
- Pacherie E (2008) The phenomenology of action: a conceptual framework. *Cognition* 107:179–217
- Papineau D (2013) In the zone. *R Inst Philos Suppl* 73:175–196
- Pickering MJ, Clark A (2014) Getting ahead: forward models and their place in cognitive architecture. *Trends Cogn Sci* 18:451–456
- Putt SS, Woods AD, Franciscus RG (2014) The role of verbal interaction during experimental bifacial stone tool manufacture. *Lithic Technology* 39(2):96–112
- Putt SS, Wijekumar S, Franciscus RG, Spencer JP (2017) The functional brain networks that underlie Early Stone Age tool manufacture. *Nat Hum Behav* 1:0102
- Railton P (2006) Normative guidance. In: Shafer-Landau R (ed) *Oxford studies in meta-ethics*, vol 1. Oxford University Press, Oxford, pp 3–34
- Railton P (2014) The affective dog and its rational tale: intuition and attunement. *Ethics* 124(4):813–859
- Railton P (2017a) Moral learning: conceptual foundations and normative relevance. *Cognition* 167:172–190
- Railton P (2017b) At the core of our capacity to act for a reason: the affective system and evaluative model-based learning and control. *Emot Rev* 9(4):335–342
- Rietveld E (2008) Situated normativity: the normative aspect of embodied cognition in unreflective action. *Mind* 117:973–1001
- Rudolf von Rohr C, Burkart JM, van Schaik CP (2011) Evolutionary precursors of social norms in chimpanzees: a new approach. *Biol Philos* 26:1–30
- Rudolf von Rohr C, van Schaik CP, Kissling A, Burkart JM (2015) Chimpanzees' bystander reactions to infanticide: an evolutionary precursor of social norms? *Hum Nat* 26:143–160
- Schlingoff L, Moore R (2017) Do chimpanzees conform to social norms? In: Andrews K, Beck J (eds) *The Routledge handbook of the philosophy of animal minds*. Routledge, London, pp 381–389
- Schmidt RA, Wrisberg CA (2008) *Motor learning and performance: a situation-based learning approach*, 4th edn. Human Kinetics, Champaign, IL
- Shennan S (2002) *Genes, memes and cultural history: Darwinian archaeology and cultural evolution*. Thames and Hudson, London
- Shennan SJ, Steele J (1999) Cultural learning in hominids: a behavioural ecological approach. In: Box HO, Gibson KR (eds) *Mammalian social learning: comparative and ecological perspectives*. Cambridge University Press, Cambridge
- Shipton C (2010) Imitation and shared intentionality in the Acheulean. *Cambridge Archaeol J* 20:197–210
- Shipton C (2019) The evolution of social transmission in the Acheulean. In: Overmann KA, Coolidge FL (eds) *Squeezing minds from stones: cognitive archaeology and the evolution of the human mind*. Oxford University Press, New York, pp 332–354
- Shipton C, Nielsen M (2015) Before cumulative culture: the evolutionary origins of overimitation and shared intentionality. *Hum Nat* 26:331–345
- Shipton C, White M (2020) Handaxe types, colonization waves, and social norms in the British Acheulean. *J Archaeol Sci Rep* 31:102352
- Shipton C, Nielsen M, Di Vincenzo F (in press). The Acheulean origins of normativity
- Sripada CS (2005) Punishment and the strategic structure of moral systems. *Biol Philos* 20:767–789

- Sripada CS, Stich S (2006) A framework for the psychology of norms. In: Carruthers C, Laurence S, Stich S (eds) *The innate mind, vol 2. culture and cognition*. Oxford University Press, New York, pp 280–301
- Stanley J, Krakauer JW (2013) Motor skill depends on knowledge of facts. *Front Hum Neurosci* 7:503
- Stanley J, Williamson T (2017) Skill. *Noûs* 51:714–726
- Sterelny K (2012a) *The evolved apprentice: how evolution made humans unique*. MIT Press, Cambridge, MA
- Sterelny K (2012b) Language, gesture, skill: the co-evolutionary foundations of language. *Philos Trans R Soc Lond B Biol Sci* 367:2141–2151
- Sterelny K (2012c) Morality's dark past. *Analyse & Kritik* 2012:95–115
- Sterelny K (2014) A Paleolithic reciprocation crisis: symbols, signals, and norms. *Biol Theory* 9:65–77
- Stout D (2010) The evolution of cognitive control. *Top Cogn Sci* 2:614–630
- Stout D (2011) Stone toolmaking and the evolution of human culture and cognition. *Philos Trans R Soc Lond B Biol Sci* 366:1050–1059
- Stout D, Chaminade T (2012) Stone tools, language and the brain in human evolution. *Philos Trans R Soc Lond B Biol Sci* 367:75–87
- Stout D, Toth N, Schick K, Chaminade T (2008) Neural correlates of Early Stone Age toolmaking: technology, language and cognition in human evolution. *Philos Trans R Soc Lond B Biol Sci* 363:1939–1949
- Stout D, Hecht E, Khreisheh N, Bradley B, Chaminade T (2015) Cognitive demands of Lower Paleolithic toolmaking. *PLoS ONE* 10(4):e0121804
- Sutton RE (2004) Emotional regulation goals and strategies of teachers. *Soc Psychol Educ* 7:379–398
- Sutton J (2007) Batting, habit and memory: the embodied mind and the nature of skill. *Sport Soc* 10:763–786
- Sutton J (2013) Skill and collaboration in the evolution of human cognition. *Biological Theory* 8:28–36
- Sznycer D, Tooby J, Cosmides L, Porat R, Shalvi S, Halperin E (2016) Shame closely tracks the threat of devaluation by others, even across cultures. *Proc Natl Acad Sci USA* 113:2625–2630
- Tangney JP, Stuewig J, Malouf ET, Youman K (2013) Communicative functions of shame and guilt. In: Sterelny K, Joyce R, Calcott B, Fraser B (eds) *Cooperation and its Evolution*. MIT Press, Cambridge, MA, pp 485–502
- Tennie C, Premo LS, Braun DR, McPherron SP (2017) Early stone tools and cultural transmission: resetting the null hypothesis. *Curr Anthropol* 58:652–654
- Tomasello M (2016) *A natural history of human morality*. Harvard University Press, Cambridge, MA
- Tomasello M (2020) The moral psychology of obligation. *Behav Brain Sci* 43:E56. <https://doi.org/10.1017/S0140525X19001742>
- Toner J, Montero BG, Moran A (2015) Considering the role of cognitive control in expert performance. *Phenomenol Cogn Sci* 14:1127–1144
- van Leeuwen EJC, Cronin KA, Haun DBM, Mundry R, Bodamer MD (2012) Neighbouring chimpanzee communities show different preferences in social grooming behaviour. *Proc R Soc Lond B Biol Sci* 279:4362–4367
- van Leeuwen EJC, Cronin KA, Haun DBM (2014) A group-specific arbitrary tradition in chimpanzees (*Pan troglodytes*). *Anim Cogn* 17:1421–1425
- Wittgenstein L (1966) *Lectures and conversations on aesthetics, psychology and religious belief*, ed Barrett C. Blackwell, Oxford
- Wolpert DM, Ghahramani Z (2000) Computational principles of movement neuroscience. *Nat Neurosci* 3:1212–1217
- Wolpert DM, Kawato M (1998) Multiple paired forward and inverse models for motor control. *Neural Networks* 11:1317–1329
- Wynn T (1993) Two developments in the mind of early Homo. *J Anthropol Archaeol* 12:299–322
- Wynn T, McGrew WC (1989) An ape's view of the Oldowan. *Man* 24:383–398

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.