

Preprint, 14 April 2020

## **Toolmaking and the Origin of Normative Cognition**

**Jonathan Birch**

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

[j.birch2@lse.ac.uk](mailto:j.birch2@lse.ac.uk)

<http://personal.lse.ac.uk/birchj1>

11,947 words

**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 skills related to toolmaking and tool use, are guided by internally represented norms of correct performance. Moreover, it is plausible that core components of human normative cognition evolved in response to the distinctive demands of transmitting complex motor skills and craft skills, especially skills related to toolmaking and tool use, through social learning. If this is correct, the expansion of the normative domain beyond technique to encompass more abstract norms of reciprocity, ritual, kinship and fairness 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.

**Key words:** normative cognition, skill, cognitive control, norms, evolution

## 1. 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.

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 capacities (to anticipate and notice norm violations, to feel moved to restore compliance, and to know what the situation requires) are core components of human normative cognition. They present 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 skills related to toolmaking and tool use, 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:

*Some core components of human normative cognition (in particular, the capacity to internally represent action-guiding norms) evolved in response to the distinctive demands of transmitting complex motor skills and craft skills, especially skills related to toolmaking and tool use, through social learning.*

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 a plausible “how-possibly” 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 (as defended by Sterelny 2012a). 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. My 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.<sup>1</sup>

## **2. 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.<sup>2</sup> The commitments of such a view are not easy to tease 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. This seems undeniable, 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

---

<sup>1</sup> Precursors to the skill hypothesis include Churchland (1996, 2000); Clark (2000) and Sterelny (2012a).

<sup>2</sup> See Fitts and Posner (1967); Dreyfus and Dreyfus (1986); Dreyfus (2005); Schmidt and Wrisberg (2008); Di Nucci (2013).

cognitively impenetrable.<sup>3</sup> This is in tension with the skill hypothesis. However, there is no good reason to believe this strong version of the “mindless” view.

The notion of expert-level skill execution as cognition-free has recently come under sustained critical pressure.<sup>4</sup> 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, who is 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 a wholly automatic affair. 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

---

<sup>3</sup> Even if there were no role for cognitive control in the execution of an *expert-level* skill, there might still be a role for cognitive control in the early stages of skill acquisition. So, it would be possible to recast the skill hypothesis as a claim about the role of normative cognition in the early stages of skill acquisition.

<sup>4</sup> See Annas (2011); Christensen et al. (2014, 2015, 2016); Montero (2016); Pacherie (2008); Papineau (2013); Pavese (2015, 2016a, b); Rietveld (2008); Sutton (2007, 2013); Stanley (2011); Stanley and Krakauer (2013); Stanley and Williamson (2001, 2017).

motor skills and craft skills—as possessing a substantial role for cognitive control. I will argue 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? As a starting point, consider Christensen and colleagues' (2015a) account of 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*.<sup>5</sup> 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’ proposal 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.<sup>6</sup> For example, when one grasps

---

<sup>5</sup> 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). My claims in this section concern model-based *action control* rather than model-based *learning*, but there is consilience here: a model-based account of skill helps explain how the evaluative models posited by these accounts of norm-learning (on independent grounds) become integrated with fluent agency.

<sup>6</sup> See Wolpert and Kawato (1998), Kawato (1999), Wolpert and Ghahramani (2000), Frith (2012); Blakemore et al. (2000, 2001); Franklin and Wolpert (2011); Pickering and Clark (2014).



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.<sup>7</sup>

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.<sup>8</sup> But 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

---

<sup>7</sup> The 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).

<sup>8</sup> 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.

*this* course, such as the softness of the suspension or the stopping power of the brakes (Christensen et al. 2015a, pp. 347-8). 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.

This picture of complex motor skill as action guided by a well-calibrated cognitive control model is summarised in the following initial account:<sup>9</sup>

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.

### 3. 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,

---

<sup>9</sup> I exclude from the scope of the account purely mental actions, such as mental arithmetic, and basic motor actions that require no sophisticated cognitive control or situational awareness, such as clenching one's fist.

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.<sup>10</sup> Wittgenstein, in the *Lectures on Aesthetics* (1966), urges us to move away from visualizing the paradigm 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.

---

<sup>10</sup> See also Buskell (2015).

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 nevertheless 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 success 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.

#### 4. Testing the Role of Norms in Skill Execution

What sort of experiment would test 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 far from conclusive.

First, let us consider what sort of experiment would support the model-based view of skill execution presented in Section 2. The main foil for the model-based view is the strong version of the “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 (i.e. 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 automated, 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, by contrast, 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.

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 predicts 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.

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 failures of skill, in themselves or in others (e.g. strong feelings of shame or anger), should also feel stronger affective responses to social norm violations in themselves or in others.

## **5. The Evolution of Normative Cognition**

So far, I have hypothesized a neglected link between normative cognition and model-based cognitive control of complex motor skill and craft skill. The link merits further exploration, but I hope I have done enough to motivate the first part of the skill hypothesis. I will now turn to the second part. 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.



## 5.1 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 hominins<sup>11</sup> 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. 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 there is no need to for the hominins themselves to have cognitive control models of the situation in which these norms are encoded. All they need to do is recognize rank and condition behaviour on it. Model-free associative learning would be enough to learn a conditioned fear response to an angry alpha.

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

---

<sup>11</sup> The term “hominin” encompasses the *Homo* genus and the australopithecines; in other words, modern humans and all ancestral species following the split from *Pan*.

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.

In addition to social emotions, rank sensitivity and simple cultural traditions, there were basic capacities for skill learning (including social learning of skills) and skill execution. 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 (Wolpert et al. 1998). But did the last common ancestor also possess cognitive control models, modelling the execution of a complex skill and anticipating the demands imposed by this particular situation? What is the behavioural marker of this more sophisticated form of action control?

It helps here to return to our focal example of elite mountain biking: in this case, 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). 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 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.

## 5.2 The elaboration of cognitive control

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 master-apprentice learning. Early hominins, like modern-day apprentices, acquired skills from the preceding generation through a process of guided trial and error: juveniles would attempt to recreate adult activities, such as foraging and toolmaking, while adults would actively correct their mistakes and provide a structured programme of progressively more demanding challenges. These scaffolded learning environments were complemented by cognitive and life-history adaptations: adults evolved to be tolerant and supportive of apprentices, who evolved to learn by imitation from their role 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 were extraordinary. An important milestone is Acheulean lithic technology. This technology, 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).<sup>12</sup>

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. I hypothesize that, like skilled mountain biking, the craft of flint knapping is (and was, from at least the Acheulean onwards) guided 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.<sup>13,14</sup>

---

<sup>12</sup> 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). Those sceptical about my claims regarding the Acheulean 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.

<sup>13</sup> 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.

<sup>14</sup> See Stout et al. (2015) and Putt et al. (2017) for neuroscientific evidence that Acheulean tool production engages mechanisms of cognitive control. Stout and collaborators have long argued for the importance of the Acheulean in human cognitive evolution (Stout et al. 2008; Stout 2011; Stout and Chaminade 2012).




FIGURE 1: Acheulean bifaces (drawing from the *Victoria History of the County of Kent, Volume I*, ed. William Page, 1908, London: Archibald Constable and Company, p. 312.)

### 5.3 The standardization of technique

This is evidence, then, that the *cognitive* dimension of skill as seen in modern humans was in place by the Acheulean period.<sup>15</sup> Is it also evidence of the *affective* dimension? I have argued

---

<sup>15</sup> I don't rule 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 and group-wide standardization (Ambrose 2001; Wynn and McGrew 1989). But see Morgan et al. (2015) for arguments that the received wisdom about the Oldowan may have underestimated its complexity.

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). 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. Ceri Shipton and Mark 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.<sup>16</sup> There is archaeological

---

<sup>16</sup> Tomasello's (2016) suggestion that the first norms were "role ideals" in joint action is probably compatible with my claims here, if we allow that toolmaking was an instance of a joint action. Shipton and Nielsen (2015)

evidence that techniques were indeed standardized at particular sites, with some variation between sites (Wynn 1993; Shipton 2010). These standardized techniques were passed down the generations with high fidelity. 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*.

Standardization of technique is possible when agents internalize a particular style of skill execution, in the sense of feeling negative affective responses to small departures from a cognitive control model. 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.

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 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, selection would have favoured agents *intrinsically motivated to master skills*—agents who felt satisfaction in achieving excellence in skill execution, and discontent at anything short of excellence.

---

discuss Acheulean tool manufacture as an early and important example of human collaboration, perhaps one involving joint intentionality in Tomasello’s sense.



Selection would have favoured mechanisms for targeting negative affect at those specific aspects of performance that required improvement.

#### **5.4 Social enrichment, teaching and a chicken-and-egg problem**

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 been social norms in a more demanding sense, which also connotes third-party monitoring, expectation and/or enforcement (Bicchieri 2005). These elaborations were, I suggest, connected to teaching.

Hiscock (2014) argues that, given the risk of injury inherent in even very small errors, Acheulean tool manufacture already required intentional teaching. Effective 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 foresee 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.

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, control models were co-opted for the regulation of the behaviour of other agents.

Admittedly, 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.

My response is that 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 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 1998; Shennan 2002; Mameli 2008). Prosocial dispositions towards campmates and kin would have been enough to stabilize simple forms of teaching, based on the demonstration of skills, prior to the enforcement of pedagogical norms.

### 5.5 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 (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.

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. 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.

### **5.6 The expansion of the normative domain**

All of this, I hypothesize, was there in the Acheulean period, prior to the origin of lexically rich languages, cumulative cultural evolution, or behaviourally modern practices such as ritual or artistic creativity, all of which plausibly require a pre-existing capacity for norm-guided skilled action. However, normative cognition at this stage 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 pervasiveness of social norms in micromanaging all aspects of social conduct brought with it norms of kinship, norms of ritual and religion, and much more demanding norms of fairness. I suggest that these norms, despite their apparently abstract and general character, may have been an elaboration of a basic capacity for norm-guided skill execution.

How so? Sterelny (2014) has argued that the expansion of the normative domain was driven by the "Palaeolithic reciprocation crisis": a package of coordination problems resulting from the demands of reciprocity in expanding social groups. 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. Meanwhile, simple norms of fairness plausibly arose in the context of collaborative hunting: skilful execution of a hunt would flow into skilful division of the spoils, guided by norms of equitable division that benefited the agent, in the long run, by showing them to be a trustworthy cooperation partner (Baumard 2016).

Larger groups also faced the problem of creating group cohesion by means other than one-on-one bonding (Dunbar 2014). The solution was ritual: skilful collective performances, high in emotional and mnemonic resonance. It is easy to see how norms of ritual might have begun as norms of skill execution: norms of how to dance, or make music, our way. Moreover, at some of these rituals, inter-band monogamous pairings would be arranged, guided by norms of kinship: norms of who can pair 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 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.

## 6. 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 such pressure. In sum, this is because Acheulean and post-Acheulean toolmaking is a complex skill with multiple possible styles of execution, calling for (i) sophisticated mechanisms of model-based cognitive control, (ii) group-wide standardization of technique through technical norms, and (iii) years of sustained practice, driven by an intrinsic motivation to achieve an internalized standard of correct performance. The demands of standardization and practice selected for cognitive changes that made them easier, enabling yet higher levels of standardization and yet more demanding practice regimes. Normative cognition was the adaptation that facilitated the standardization and practice of our ways of performing complex practical skills. This is an example of a “gene-culture co-evolution” hypothesis (Boyd and Richerson 1985).

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, but there is disagreement about the *type* of cooperation that mattered most.

## **6.1 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. I see an important role for Acheulean toolmaking because it seems, to me, 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. The article could just as well have been called “Complex Motor/Craft Skill and the Origin of Normative Cognition”. 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) and Baumard (2010) have also suggested, in the context of skilful mutualistic cooperation. But it was not the mutualistic or cooperative character of the behaviour that generated the selection pressure for technical norms: it was its skilful character.

## **6.2 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. 103-4). Altruism-based accounts of the evolution of normative cognition thus face the question: *If the problem was a need to boost altruism, why did natural selection not just enhance 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 mechanism. The skill hypothesis shows there were design problems—the



need to standardize technique across groups and to motivate sustained practice—to which the distinctive features of normative cognition provide a solution.

### **6.3 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. For example, Kitcher writes:

Suppose a type of altruism failure, keeping food items for oneself, say, regularly elicits aggressive retaliation from others. Chimpanzees and hominids could recognize the regularity, thus allowing for variants who recognize the potential threats to them if they fail to share, and whose fear generates compliance. (Kitcher 2011, p. 90)

For Kitcher (2011), the punishment hypothesis and the “altruism failures” hypothesis are complementary: altruism failures have adverse fitness consequences because they elicit punishment. 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. We should again note that natural selection sometimes assembles overcomplicated solutions to design problems. But again, it is an advantage of the skill hypothesis that it does not involve the same kind of mismatch between the features of normative cognition and the design problem it initially solved.

#### **6.4 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).<sup>17</sup> The same goes for teaching and the group-wide standardization of technique: these things may require signalling, but the signalling may be entirely gestural. The story is compatible with an origin for language long after the basic elements of normative cognition were in place. 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, like accounts that assume sophisticated punishment practices, face a chicken-and-egg problem.

---

<sup>17</sup> Some have, however, suggested that skilled toolmaking might have enabled and driven the evolution of language (e.g. Stout and Chaminade 2012; Sterelny 2012b).

### **6.5 Is sociality necessary?**

A final question the skill hypothesis brings into view is whether normative cognition could only ever have evolved in a richly cooperative species, or whether its evolution was only contingently entangled with the evolution of cooperation. Consider an intelligent but non-social animal, such as an octopus. Could an octopus evolve normative cognition?

The received wisdom, I take it, is that it could not: normative cognition is a response to distinctive demands imposed by social life. The punishment hypothesis and the “altruism failures” hypothesis endorse this claim. Yet the skill hypothesis suggests it may be too quick. A non-social but intelligent creature might still stand to gain from learning the right way to manufacture a complex tool. If the process takes months or years to master, there might also be selection for an intrinsic motivation to practice and improve in relation to an internalized standard of correct performance. If complex tool manufacture becomes ecologically central to this creature’s way of life, we would expect to see cognitive changes evolving in response, and these cognitive changes might resemble human normative cognition in some respects, while differing in others—since the groupwide standardization of technique would not be an issue. I am not suggesting this has happened, only that it could happen. The evolution of normative cognition might not necessarily require sociality, even though the relationship in our own case is close.

## **7. Looking for Archaeological Evidence**

In Section 3, 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 (and especially norms of toolmaking) should not post-date more abstract norms of ritual, reciprocity, fairness and kinship. It would count against the hypothesis if there are examples of societies with very simple technical norms but highly complex, highly abstract social norms. The prediction is that the complexity of the technical and non-technical norms a society can support should be correlated. 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.

In sum, although it has yet to receive serious empirical attention, I have made a case that the skill hypothesis is consistent with 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 and evolutionary anthropology.

### **Acknowledgements**

I thank Ali Boyle, Richard Bradley, Campbell Brown, Susanne Burri, Andrew Buskell, Wayne Christensen, Ellen Fridland, Mattia Gallotti, Balasz Gyenis, Cecilia Heyes, Laurenz Hudetz, Liam Kofi Bright, Peter Railton, Kim Sterelny, John Sutton, Johanna Thoma, Felix Warneken and Joeri Witteveen 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. This work was supported by a Philip Leverhulme Prize from the Leverhulme Trust and visiting fellowships at the ANU in 2015 and 2018.

## References

- Allen, Nicholas J., Hilary Callan, Robin Dunbar and Wendy James (eds). (2011). *Early Human Kinship: From Sex to Social Reproduction*. Hoboken, NJ: Wiley-Blackwell.
- Ambrose, Stanley H. (2001). Paleolithic technology and human evolution. *Science*, 291, 1748–1753.
- Andrews, Kristin. (2009). Understanding norms without a theory of mind. *Inquiry*, 52, 433–448.
- Annas, Julia. (2011). *Intelligent Virtue*. Oxford: Oxford University Press.
- Baumard, Nicolas. (2016). *The Origins of Fairness: How Evolution Explains Our Moral Nature*, trans. Paul Reeve. New York: Oxford University Press.
- Bicchieri, Cristina. (2005). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.
- Blakemore, Sarah-Jayne, Daniel Wolpert and Chris Frith. (2000). Why can't you tickle yourself? *Neuroreport*, 11, R11-R16.
- Blakemore, Sarah-Jayne, Daniel Wolpert and Chris Frith. (2000). The cerebellum is involved in predicting the sensory consequences of action. *Neuroreport*, 12, 1879-1884.
- Boehm, Christopher. (2012). *Moral Origins: The Evolution of Virtue, Altruism and Shame*. New York: Basic Books.
- Boyd, Robert and Richerson, Peter J. (1985) *Culture and the Evolutionary Process*. Chicago, IL: University of Chicago Press.

- Buskell, Andrew. (2015). How to be skilful: normative sensitivity and opportunistic robustness. *Synthese*, 192, 1445-1466.
- Chapais, Bernard. (2008). *Primeval Kinship: How Pair-Bonding Gave Birth to Human Society*. Cambridge, MA: Harvard University Press.
- Christensen, Wayne, Kath Bicknell, Doris McIlwain, and John Sutton. (2015a). The sense of agency and its role in strategic control for expert mountain bikers. *Psychology of Consciousness: Theory, Research, and Practice*, 2, 340-353.
- Christensen, Wayne, John Sutton and Doris J. F. McIlwain. (2015b). Putting pressure on theories of choking: towards an expanded perspective on breakdown in skilled performance. *Phenomenology and the Cognitive Sciences*, 14, 253-293.
- Christensen, Wayne, John Sutton and Doris J. F. McIlwain. (2016). Cognition in skilled action: meshed control and the varieties of skill experience. *Mind and Language*, 31, 37-66.
- Churchland, Paul M. (1996). The neural representation of the social world. In Larry May, Marilyn Friedman and Andy Clark (eds), *Mind and Morals: Essays on Cognitive Science and Ethics*. Cambridge, MA: MIT Press, pp. 91-108.
- Churchland, Paul M. (2000). Rules, know-how and the future of moral cognition. *Canadian Journal of Philosophy*, 30, 291-206.
- Clark, Andy. (2000). Word and action: reconciling rules and know-how in moral cognition. *Canadian Journal of Philosophy*, 30, 267-289.
- Clark, Andy. (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. New York: Oxford University Press.
- Conant, Roger C. and W. Ross Ashby. (1970). Every good regulator of a system must be a model of that system. *International Journal of Systems Science*, 1, 89-97.

- Corbey, Raymond, Adam Jagich, Krist Vaesen, and Mark Collard. (2016). The Acheulean handaxe: more like a bird's song than a Beatles' tune? *Evolutionary Anthropology*, 25, 6-19.
- Crockett, Molly J. (2013). Models of morality. *Trends in Cognitive Sciences*, 17(8), 363-366.
- Cushman, Fiery. (2013). Action, outcome, and value a dual-system framework for morality. *Personality and Social Psychology Review*, 17(3), 273-292.
- Cushman, Fiery, Victor Kumar and Peter Railton. (2017). Moral learning: psychological and philosophical perspectives. *Cognition*, 167, 1-10.
- Davidson, Ian. (2002). The 'finished artefact fallacy': Acheulean handaxes and language origins. In Alison Wray (ed.) *The Transition to Language*. Oxford: Oxford University Press, pp. 180-203.
- Di Nucci, Ezio. (2013). *Mindlessness*. Cambridge: Cambridge Scholars Publishing.
- Dreyfus, Hubert L. (2005). Overcoming the myth of the mental: how philosophers can profit from the phenomenology of everyday expertise. *Proceedings and Addresses of the American Philosophical Association*, 79, 47-65.
- Dreyfus, Hubert L. and Stuart E. Dreyfus. (1986). *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. New York: Free Press.
- Dreyfus, Hubert L. and Stuart E. Dreyfus. (1991). Towards a phenomenology of ethical expertise. *Human Studies*, 14, 229-250.
- Dunbar, Robin. (2014). *Human Evolution*. London: Pelican.
- Fessler, Daniel M. T. (1999). Toward an understanding of the universality of second order emotions. In Alexander Laban Hilton (ed.), *Beyond Nature or Nurture: Biocultural Approaches to the Emotions*. Cambridge: Cambridge University Press, pp. 75-116.
- Fessler, Daniel M. T. (2004). Shame in two cultures: implications for evolutionary approaches. *Journal of Cognition and Culture*, 4, 207-262.

- Fessler, Daniel M. T. (2007). From appeasement to conformity: evolutionary and cultural perspectives on shame, competition, and cooperation. In Jessica L. Tracy, Richard W. Robins and June Price Tangney (eds), *The Self-Conscious Emotions: Theory and Research*. New York: Guildford Press, pp. 174-193.
- Fitts, Paul Morris and Posner, Michael I. (1967). *Human Performance*. Belmont, CA: Brooks/Cole Publishing.
- Franklin, David W. and Daniel M. Wolpert. (2011). Computational mechanisms of sensorimotor control. *Neuron*, 72, 425-442.
- Fridland, Ellen. (2017). Automatically minded. *Synthese*, 194, 4337-4363.
- Frith, Chris. (2012). Explaining delusions of control: the comparator model 20 years on. *Consciousness and Cognition*, 21, 52-54.
- Gibbard, Allan. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Graziano, Paulo A., Rachael D. Reavis, Susan P. Keane, and Susan D. Calkins. (2007). The role of emotion regulation and children's academic success. *Journal of School Psychology*, 45, 3-19.
- Greene, Joshua D. (2017). The rat-a-gorical imperative: moral intuition and the limits of affective learning. *Cognition*, 167, 66-77.
- Gullone, Eleanora, Elizabeth K. Hughes, Neville J. King, and Bruce Tonge. (2010). The normative development of emotion regulation strategy use in children and adolescents: a 2-year follow-up study. *Journal of Child Psychology and Psychiatry*, 51, 567-574.
- Haidt, Jonathan. (2001). The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological Review*, 108, 814-834.
- Hewlett, Barry S. and L. L. Cavalli-Sforza. (1986). Cultural transmission among Aka pygmies. *American Anthropologist*, 88, 922-934.



- Hiscock, Peter. (2014). Learning in lithic landscapes: a reconsideration of the hominid “toolmaking” niche. *Biological Theory*, 9, 27-41.
- Humle, Tatyana. (2010). How are army ants shedding new light on culture in chimpanzees? In Elizabeth V. Lonsdorff, Stephen R. Ross and Tetsuro Matsuzawa (eds), *The Mind of the Chimpanzee: Ecological and Experimental Perspectives*. Chicago, IL: University of Chicago Press, pp. 116-127.
- Iovita, Radu, Inbal Tuvi-Arad, Marie-Hélène Moncel, Jackie Despriée, Pierre Voinchet, Jean-Jacques Bahain. (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, Richard. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Kawato, Mitsuo. (1999). Internal models for motor control and trajectory planning. *Current Opinion in Neurobiology*, 9, 718-727.
- Kitcher, Philip. (2011). *The Ethical Project*. Cambridge, MA: Harvard University Press.
- Luncz, Lydia V., Roger Mundry and Christopher Boesch. (2012). Evidence for cultural differences between neighboring chimpanzee communities. *Current Biology*, 22, 922-926.
- Lycett, Stephen J. (2008). Acheulean variation and selection: does handaxe symmetry fit neutral expectations? *Journal of Archaeological Science*, 35(9), 2640-2648.
- Machery, Eduoard and Ron Mallon. (2010). Evolution of morality. In John Doris (ed.), *The Moral Psychology Handbook*. Oxford: Oxford University Press, pp. 3-47.
- Maibom, Heidi L. (2010). The descent of shame. *Philosophy and Phenomenological Research*, 80(3), 566-594.
- Mameli, Matteo. (2008). Understanding culture: A commentary on Richerson and Boyd’s Not By Genes Alone. *Biology and Philosophy*, 23, 269-281.

- Montero, Barbara Gail. (2006). Proprioception as an aesthetic sense. *Journal of Aesthetics and Art Criticism*, 64, 231-242.
- Montero, Barbara Gail. (2016). *Thought in Action: Expertise and the Conscious Mind*. Oxford: Oxford University Press.
- Morgan, Thomas J. H., Natalia T. Uomini, Luke E. Rendell, Laura Chouinard- Thuly, Sharon E. Street, Hannah M. Lewis, Catherine P. Cross, et al. (2015). Experimental evidence for the co-evolution of hominin tool-making teaching and language. *Nature Communications*, 6, 1-8.
- Murphy, Michael and Rhea A. White. (1995). *In the Zone: Transcendent Experience in Sports*. New York: Penguin.
- Pacherie, Elisabeth. (2008). The phenomenology of action: a conceptual framework. *Cognition*, 107, 179-217.
- Papineau, David. (2013). In the zone. *Royal Institute of Philosophy Supplement*, 73, 175-196.
- Pavese, Carlotta. (2015). Knowing a rule. *Philosophical Issues*, 25, 165-188.
- Pavese, Carlotta. (2016a). Skill in epistemology I: Skill and knowledge. *Philosophy Compass*, 11, 642-649.
- Pavese, Carlotta. (2016b). Skill in epistemology II: Skill and know how. *Philosophy Compass*, 11, 650-660.
- Pickering, Martin J. and Andy Clark. (2014). Getting ahead: forward models and their place in cognitive architecture. *Trends in Cognitive Sciences*, 18, 451-456.
- Putt, Shelby S., Alexander D. Woods, and Robert G. Franciscus. (2014). The role of verbal interaction during experimental bifacial stone tool manufacture. *Lithic Technology*, 39(2), 96–112.

- Putt, Shelby S., Sobanawartiny Wijekumar, Robert G. Franciscus and John P. Spencer. (2017). The functional brain networks that underlie Early Stone Age tool manufacture. *Nature Human Behaviour*, 1, 0102.
- Railton, Peter. (2006). Normative guidance. In Russ Shafer-Landau (ed.), *Oxford Studies in Meta-Ethics, Vol. 1*. Oxford: Oxford University Press, pp. 3-34.
- Railton, Peter. (2014). The affective dog and its rational tale: intuition and attunement. *Ethics*, 124(4), 813-859.
- Railton, Peter. (2017a). Moral learning: conceptual foundations and normative relevance. *Cognition*, 167, 172-190.
- Railton, Peter. (2017b). At the core of our capacity to act for a reason: the affective system and evaluative model-based learning and control. *Emotion Review*, 9(4), 335-342.
- Rietveld, Erik. (2008). Situated normativity: the normative aspect of embodied cognition in unreflective action. *Mind*, 117, 973-1001.
- Schlingoff, Laura and Richard Moore. (2017). Do chimpanzees conform to social norms? In Kristin Andrews and Jacob Beck (eds), *The Routledge Handbook of the Philosophy of Animal Minds*. London: Routledge, pp. 381-389.
- Schmidt, Richard A. and Wrisberg, Craig A. (2008). *Motor Learning and Performance: A Situation-Based Learning Approach*. 4<sup>th</sup> edition. Champaign, IL: Human Kinetics.
- Shennan, Stephen (2002). *Genes, Memes and Cultural History: Darwinian Archaeology and Cultural Evolution*. London: Thames and Hudson.
- Shennan, Stephen J. and Steele James (1999). Cultural learning in hominids: a behavioural ecological approach. In Hilary O. Box and Kathleen R. Gibson (eds.), *Mammalian Social Learning: Comparative and Ecological Perspectives*. Cambridge: Cambridge University Press.

- Shipton, Ceri. (2010). Imitation and shared intentionality in the Acheulean. *Cambridge Archaeological Journal*, 20, 197-210.
- Shipton, Ceri and Nielsen, Mark. (2015). Before cumulative culture: the evolutionary origins of overimitation and shared intentionality. *Human Nature*, 26, 331-345.
- Sripada, Chandra Sekhar (2005). Punishment and the Strategic Structure of Moral Systems. *Biology & Philosophy*, 20, 767–789.
- Sripada, Chandra Sekhar and Stephen Stich. (2006). A framework for the psychology of norms. In Peter Carruthers, Stephen Laurence and Stephen Stich (eds), *The Innate Mind Volume 2: Culture and Cognition*. New York: Oxford University Press, pp. 280-301.
- Stanley, Jason. (2011). *Know How*. Oxford: Oxford University Press.
- Stanley, Jason and John W. Krakauer. (2013). Motor skill depends on knowledge of facts. *Frontiers in Human Neuroscience*, 7, 503.
- Stanley, Jason and Timothy Williamson. (2001). Knowing how. *Journal of Philosophy*, 98, 411-444.
- Stanley, Jason and Timothy Williamson. (2017). Skill. *Noûs*, 51, 714-726.
- Sterelny, Kim. (2012a). *The Evolved Apprentice: How Evolution Made Humans Unique*. Cambridge, MA: MIT Press.
- Sterelny, Kim. (2012b). Language, gesture, skill: the co-evolutionary foundations of language. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 367, 2141-2151.
- Sterelny, Kim. (2012c). Morality's dark past. *Analyse & Kritik*, 2012, 95-115.
- Sterelny, Kim. (2014). A Paleolithic reciprocation crisis: symbols, signals, and norms. *Biological Theory*, 9, 65-77.

- Stout, Dietrich. (2011). Stone toolmaking and the evolution of human culture and cognition. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 366, 1050-1059.
- Stout, Dietrich and Thierry Chaminade. (2012). Stone tools, language and the brain in human evolution. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 367, 75-87.
- Stout, Dietrich, Nicholas Toth, Kathy Schick and Thierry Chaminade. (2008). Neural correlates of Early Stone Age toolmaking: technology, language and cognition in human evolution. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 363, 1939-1949.
- Stout, Dietrich, Erin Hecht, Nada Khreisheh, Bruce Bradley and Thierry Chaminade. (2015). Cognitive demands of Lower Paleolithic toolmaking. *PLoS ONE*, 10(4), e0121804.
- Sutton, John. (2007). Batting, habit and memory: the embodied mind and the nature of skill. *Sport in Society: Cultures, Commerce, Media, Politics*, 10, 763-786.
- Sutton, John. (2013). Skill and collaboration in the evolution of human cognition. *Biological Theory*, 8, 28-36.
- Sutton, Rosemary E. (2004). Emotional regulation goals and strategies of teachers. *Social Psychology of Education*, 7, 379-398.
- Sznycer, Daniel, John Tooby, Leda Cosmides, Rodi Porat, Shaul Shalvi and Eran Halperin. (2016). Shame closely tracks the threat of devaluation by others, even across cultures. *Proceedings of the National Academy of Sciences of the United States of America*, 113, 2625-2630.
- Tangney, June P., Jeffrey Stuewig, Elizabeth T. Malouf and Kerstin Youman. (2013). Communicative functions of shame and guilt. In Kim Sterelny, Richard Joyce, Brett

- Calcott, and Ben Fraser (eds.), *Cooperation and its Evolution*. Cambridge, MA: MIT Press, pp. 485-502.
- Tennie, Claudio, L. S. Premo, David R. Braun, and Shannon P. McPherron. (2017). Early stone tools and cultural transmission: resetting the null hypothesis. *Current Anthropology*, 58, 652-654.
- Tomasello, Michael. (2014). *A Natural History of Human Thinking*. Cambridge, MA: Harvard University Press.
- Tomasello, Michael. (2016). *A Natural History of Human Morality*. Cambridge, MA: Harvard University Press.
- Toner, John, Barbara Gail Montero and Aidan Moran. (2015). Considering the role of cognitive control in expert performance. *Phenomenology and the Cognitive Sciences*, 14, 1127-1144.
- van Leeuwen, Edwin J. C., Katherine A. Cronin, Daniel B. M. Haun, Roger Mundry and Mark D. Bodamer. (2012). Neighbouring chimpanzee communities show different preferences in social grooming behaviour. *Proceedings of the Royal Society of London B: Biological Sciences*, 279, 4362-4367.
- van Leeuwen, Edwin J. C., Katherine A. Cronin, and Daniel B. M. Haun. (2014). A group-specific arbitrary tradition in chimpanzees (*Pan troglodytes*). *Animal Cognition*, 17, 1421-1425.
- Wittgenstein, Ludwig. (1966). *Lectures and Conversations on Aesthetics, Psychology and Religious Belief*, ed. Cyril Barrett. Oxford: Blackwell.
- Wolpert, Daniel M. and Zoubin Ghahramani. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, 3, 1212-1217.
- Wolpert, Daniel M. and Mitsuo Kawato. (1998). Multiple paired forward and inverse models for motor control. *Neural Networks*, 11, 1317-1329.

Wynn, Thomas. (1993). Two developments in the mind of early Homo. *Journal of Anthropological Archaeology*, 12, 299-322.

Wynn, Thomas, and W. C. McGrew. (1989). An ape's view of the Oldowan. *Man*, 24, 383–398.