Kim Sterelny*

The Skill Hypothesis: A Variant

https://doi.org/10.1515/auk-2021-0013

Abstract: The basic idea of Birch's analysis is plausible: normative guidance began in agents' assessment of their own craft skills. But I suggest developing that idea in a different way. I suggest that proto-normative affect plays its guiding role diachronically, in the development of those skills, rather than synchronically, in modulating their moment-by-moment execution. More importantly, I suggest a different pathway to normative affect's direction at second and third parties. Normative response became social in the context of skilled collaborative activities, for in those activities others' failures have material consequences for each agent. In such collaborations, all have reason to care about others' skill, or lack of it.

Keywords: origin of norms; normative guidance; norms and collaboration; norms and skills; Jonathan Birch

Jonathan Birch develops his skill hypothesis in Birch (2021), and summarises it in this issue. It is inevitably somewhat speculative. But I find the core suggestion—that norms had their evolutionary origin in the origins and expansion of expertise—very plausible. As Jonathan emphasises, many forager craft and field skills are very demanding. They develop only through sustained effort and practice. The ethnography of documented forager communities suggests that this effort is not imposed. The ethnography of foragers converges on the theme that forager children are free-range, spending much of their time in self-directed mixed age groups, learning their trade through a mixture of individual exploration, collaborative learning, horizontal social transmission (with younger children learning from older ones), eavesdropping on the adult world in the physically and socially intimate world of the overnight camp, all leavened with occasional adult teaching and advice.[1] So urgent practical necessity does not require forager children to forage, though they

1 For recent reviews of this literature, see: Boyette 2016; Lew-Levy/Reckin et al. 2017; Boyette/Hewlett 2018; Lew-Levy/Lavi et al. 2018. For a series of case studies, see Hewlett/Fouts et al. 2011; Terashima/Hewlett 2016.

***Corresponding author: Kim Sterelny,** School of Philosophy, Australian National University, Australia, e-mail: Kim.Sterelny@anu.edu.au.

often do indeed forage, sometimes with tools they make themselves, and in doing so they contribute to the family economy. But they do not typically work very hard or very long, and their craft and foraging efforts are often blended with play and socialising. Moreover, their skill acquisition is not imposed or micro-managed by adults, backed by the implicit or explicit threats of sanction or promises of reward.

This autonomous, self-directed developmental trajectory is nicely captured in the following report from the Indian subcontinent:

> Throughout three successive days, the boys [10 and 11 years old] did practically nothing other than intensive experimentation with trapping, during which they tackled various kinds of technical shortcomings in their traps. These technical problems were overcome through repeated bouts of trial-and-error learning. It is important to note that throughout that time, many adults, fully knowledgeable and experienced in trapping, came and went along a path that stretches along this brook. Occasionally, one adult or another (including Rajan's father) stopped for a brief moment and, from some distance, silently observed the boys while they set their traps or tried to fix them. None of the passers-by approached the boys, not even once, nor did the boys come to ask for any advice or help from these adults. Throughout this time the boys did not ask for any guidance. (Naveh 2016, 128)

In short: forager children do not need to be forced to learn their many skills. They find skill acquisition intrinsically rewarding. While we cannot project the ethnography of documented forager childhoods onto the deep past, the core challenges of becoming a competent mid-Pleistocene forager were probably relevantly similar to those known from ethnography. Forager subsistence even then very likely depended on information-intensive field skills: tracking, and being able to read the environment more generally (for example recognising safe and secure water sources at a distance); recognising a wide variety of plants in differing seasonal conditions; recognising safe camp sites. While we have no direct knowledge of their soft material technology, the economics of gathering probably required some form of container technology.[2] Their lithic technology is known to be demanding, at least from about 1 mya, with the emergence of more refined versions of Acheulian tools (Kuhn 2020) (and perhaps earlier, from the origins of the Acheulian).

So I agree that it is quite likely that skill acquisition in the mid-Pleistocene depended on affect: pride in good, efficient performance; pleasure in recognising your improving capacities as skill acquisition develops; dissatisfaction when your efforts fall short of your own expectations, and sometimes a determination to do

2 I assume here that mid-Pleistocene humans practiced some form of home base foraging, with different foraging parties returning to some established camp site; it would be very difficult for a group to both hunt and gather without a home base economy of some kind, and hunting without gathering was probably too risky: Wrangham 2017.

better. Pride in performance is the mother of normative guidance. I take myself to be agreeing with Jonathan here, though in developing this thought, he makes much of directed *dissatisfaction*, saying little about the importance of positive affect. My hunch is that positive affect is of at least equal importance. As I dimly recall from my few moments of triumph in cricket, there is an enormous buzz when things go right.

However, while endorsing the central element of the skill hypothesis, I propose developing it in a somewhat different way. First, I would emphasise longer rather than shorter temporal scales. In the one individual, skills are developed and deployed over temporal scales that vary from decades to milliseconds. While affect—a sense of matters going well or badly—probably plays some role in skill acquisition and execution on all those scales, in my view, proto-normative affect is important over greater time depths than those Jonathan emphasises. He agrees that satisfaction and dissatisfaction play an important role in motivating practice, and that our sensitivity to these emotions was in part selected for their role in practice. However his discussion emphasises synchronic skill execution. Second and more important, I doubt that the local standardisation of craft skills played the role he attributes to it. We both think that proto-norms evolve initially because they regulated individual crafting. Later, these affective responses took on an additional interpersonal role. Agents responded affectively not just to their own actions but to those of others. Jonathan suggests that this expansion was mediated by locally shared standards of correct craft performance. I am sceptical, suggesting instead that it was driven by collaborative foraging. In this setting, others' errors impose costs on you.

First, time-depth. An important element of the skill hypothesis is that skill is cognitively rich, even when skilled agents cannot verbally analyse the organisation of their skill. Borrowing from the work of Wayne Christiansen and John Sutton, Jonathan captures this picture of skill through the idea of a cognitive control model. Through these models, agents represent the causal organisation of the domain of action (in particular, how agent action influences a substrate). This causal model structures their choice and mode of a motor sequence. It does so even in highly time pressured action, even when that sequence is ballistic once launched: it is, for example, when an off-break bowler is about to release the ball at the point of delivery. These causal models are not module-like in the classic Fodorian sense (Fodor 1983). A model is not encapsulated or cognitively impenetrable. It is at least somewhat open to introspection, and it is typically the result of an extensive and specific learning history. Following Christiansen and Sutton, Jonathan develops this idea through the persuasive example of elite mountain bike riding. In doing so, he emphasises the importance of this control model in shaping even in-the-moment time-pressured skill execution. Without at all denying the importance of

the causal control model to even these ballistic, time-pressured motor sequences, I would place greater emphasis on the role of the model in more diachronic contexts: One is a context of 'policy choice': the agent knows that he/she is going to be exercising her skill in a particular context, and decides on an overall strategy that will influence a sequence of skilled actions. A batsman knows that a pitch is playing low, and decides on a policy of playing forward; a conversationalist, knowing she will be talking with a valued but straight-laced and conservative grandmother, chooses a less racy conversational register; a chess player knows the preferences of her opponent, and so steers the game into a more closed rather than more open pawn structure. Over a still longer time frame, the causal model shapes the agent's skill development. For the model enables the agent to diagnose errors in, and limits to, her performance. It shapes her choice of particular practice routines, and regulates (through her assessment of her skill level) the contexts in which her skill is deployed.

I would emphasise these longer time frames. In my view, it is only over these longer time frames that something akin to norms has a clear footprint. I agree there is some affectively response tied to in-the-moment time-pressured skill execution. But it is affect in the most minimal sense. As Jonathan says, in action execution where there is a mis-match between the predicted sensory flow and the actual sensory flow, the action just feels wrong. But this form of affect is found in actions we do not normally think of as the execution of a skill. For example, in walking down some rough bush steps, when my foot does not feel firm substrate when I expect it to, it feels wrong, and I adjust my step, and perhaps at the same time grab a convenient tree. This feels-wrong prompts a rapid, semi-automatic adjustment. But the affect is neither intense, lasting, nor (most importantly) is it a matter of self-reproach. Of course, if my step really fails, with the substrate crumbling under my weight, and I begin to fall dangerously, my emotional response to the gap between predicted and actual experience will be intense. But it is alarm rather than normative disgust; I am not berating myself for clumsiness as I topple over. In my view, despite this role for an affective salient mis-match between expected and actual sensory flow, that does not make walking a normatively guided skill. That is so even though it is multi-dimensional: I can stumble and stagger, triggering the that-feels-wrong response in countless different ways. I conjecture that a horse can stumble and recover, with its recovery probably triggered by the contrast between expected and actual sensory experience.

I am more inclined to doubt that horses can decide on policies for safer walking over difficult ground. We can, and over the time frames in which policies are set, successfully or not, affect is more enduring and more obviously normative. Thus I might feel self-reproach if I knew, or should have known, that the path was unstable, and failed to respond by walking more carefully, or with an alpine stick for extra

support. That is a policy failure: failing to calibrate an activity to a context. For adults, routine walking is almost completely automatic. But as walking becomes a component of activities like orienteering, canyoning or the more demanding forms of hiking, it becomes increasingly an exemplar of a paradigm skill, and in a coarse-grained way, we decide how we will walk, and that can indeed result in profound satisfaction or its opposite. Over still deeper time frames, in my chess playing days I became increasingly aggravated by persistently drifting into poor positions while playing black in queen's pawn openings, and made a systematic effort at focussed practice around the Nimzo-Indian Defence. So I am inclined to think that in the context of skill execution, for affect to be something like 'directed discontent' (or 'directed approval'), and so for it to be a form (or early version) of normative regulation, it has to do more than induce an on-the-fly change in a specific action sequence. It has to play some more enduring role. Perhaps it shapes the policy for the iterated and strategic execution of a skill. Perhaps it shapes the way the agent develops and maintains the skill.

So I see craft norms as organising action over time, as organising sequences in the light of the causal model of the specific context in which those actions will play out, and as motivating skill development through practice, and targeting that practice on specific aspects of an agent's skill. This disagreement is probably just a matter of emphasis. More seriously, I am sceptical of the role local standardisation of craft norms plays in his account. For him, this local standardisation makes an individual's standards of correct execution directed not just at his/her own performance, but at the performance of others. He takes these generalised impersonal norms to be in place at some stage in the Acheulian,[3] and they are norms in a quite full-blooded sense, for they are internally represented and intrinsically motivating standards of correct performance, though their domain is still limited. These norms regulate an individual's craft productions, leading to satisfaction when his/her own performance conforms to these standards, and dissatisfaction, not just when his own productions fall short of those standards, but likewise when others' productions fall short as well. In his view, these Acheulian norms are in-

3 Very reasonably, he does not say when in the Acheulian, and in which lineages. The Acheulian was long, and a number of presumptively different hominin species used Acheulian tools. These include *erectus, ergaster, Heidelbergensis*, early Neanderthals and probably our lineage. There is likely to have been considerable variation in culture and cognition amongst Acheulian tool-users. So perhaps the hypothesis is: towards the end of the Acheulian there were some communities where practices were uniform, and the uniformity was in part sustained by mutual monitoring of one another's performances (and it is indeed likely that tool-making was public enough for that to be possible). But given the homogenising tendencies of oblique forms of social learning, I doubt that the record supports even this cautious version of Jonathan's view.

terpersonal even though they do not yet regulate social behaviour. Rather, they play an important but limited role in the life of Acheulian hominins: they regulate communal craft activities according to locally shared standards. As he sees it, the local standardisation of lithic craft technique is the archaeological signal of shared community norms. Importantly, these are norms in something like our sense of norms. Serious failure to live up to them triggers normative anger or disgust in others; guilt or shame in the perpetrator. The emergence of these local shared standards of correctness mediates the shift from self-directed normative guidance to normatively guided attitudes to second and third persons.

I am unconvinced. First, the evidence for local uniformity is not strong (certainly not uniformly across the Acheulian; for a recent sceptical view of uniformity claims, see (Wilkins 2020)). It is true that Ceri Shipton cites few persuasive cases arguing for the social transmission of technique; that is, something like true imitation learning, rather more coarse-grained forms of social learning (Shipton 2019). But at these sites there seems to be a significant *bias* in technique choice, in a context in which several functionally equivalent options are available. They are *not* cases where everyone, or almost everyone, is using the same technique. One technique rather than another is in majority use, but that is all. Second, local convergence on technique (especially partial convergence) does not require norms. Various forms of oblique social learning will produce the convergence in method Shipton and others document (see as well Stout/Rogers et al. 2019). Conformist social learning, copy-the-best, copy-the-most prestigious can all produce local convergence to various degrees. Shared default practices are one thing; shared standards in which agents notice and respond to the artisan practices of others with approval or disapproval is another. The first can appear without the second, and so the record does not support the second.

Does collaboration make the difference? Would it require shared norms? As Ceri Shipton shows, there is some evidence of collaborative stone tool making (Shipton/Nielson 2015, and for a more recent example, see Hiscock 2004). But that is still not evidence of a shared community standard. First, even if collaboration requires the two artisans to make their tools the same way, that only shows convergent practice. Neither need be aware of, or take any notice of, the practices of others. Each simply assumes that the other is working in a similar way. Indeed, the very possibility of alternative techniques might never occur to either. Second, it is not at all obvious the collaboration does require uniformity in technique: I quarry a large core, and knap it into a cobble for you to shape. So long as I get back a cleaver with a robust cutting edge and a safe grip, I may not care how you do it. Even if my preferences are more fine-grained, collaboration can depend on mutual awareness of one another's lithic preferences, rather than a single standardised preference. Finally, while there are indeed some intriguing cases of collaboration,

there is no evidence that collaboration was the typical mode of making tools in the Acheulian in any hominin lineage.

Could teaching drive a shift to craft norms directed at, assessing and hence shaping the craft work of one's social partners? Both Jonathan and I find persuasive Peter Hiscock's argument that the high cost of learning to work stone selected for its active teaching (Hiscock 2014). But, first, teaching is perfectly compatible with different individuals in the community teaching their own preferred techniques. Indeed, there is ethnographic evidence that forager cultures differ in the extent to which technique is uniform. In Richard Gould's discussion of the stone working techniques of Australian foragers (Gould 1980), he notes that Central Desert Australian foragers (in the areas around Alice Springs) have an elaborated technical vocabulary to discuss stone and stone making techniques, with clearly shared recognition of expert performance, which is highly valued. This is a lithic culture which will tend to produce uniform technique. In contrast, Western Desert foragers have a much less rich vocabulary, use expedient tools a lot more, and there is much less evidence of cultural convergence on shared templates. Moreover, while forager ethnography supports the idea that the teaching is important, it also indicates that for the most part, the teaching of *craft skills* is unobtrusive, not directive, and not moralised.[4]

Given my scepticism about the evidence for shared craft norms in the Acheulian, I am not persuaded that all of the idea that shame and anger (normative outrage) became incorporated into Acheulian craft norms, as conformity to the locally-adopted standard became an important social signal in partner choice (an idea developed in Birch 2021, 5.5). I agree that partner choice and hence social reputation was important in Acheulian social worlds. But it is hard to see why conformity to local craft practices would be recruited as a signal of social reliability. Why would communities rely on *indirect markers* of prosociality like respect for craft norms? For these are very intimate social worlds, in which agents have long shared histories. So we should expect reputation to depend directly on the relevant aspects of an agent's history: as a good co-operator; a valuable ally; a skilled hunter; a knowledgeable gatherer; a dangerous enemy. For these were probably communities in which individual histories were known. Ethnographically known forager communities can be quite large and multi-level, with residential groups embedded in larger communities. In these larger communities, individual histories might not be known. But these complex communities probably depended on full language, hence may well be late Pleistocene (Sterelny 2019). In any case, even

4 Teaching tends to be more directive when it is teaching ritual, esoteric knowledge, and social norms themselves.

in ethnographically described forager worlds, most social interaction is within residential groups, and between agents who know one another intimately well.

I agree craft performance can be a signal. But most plausibly, when it functions as a signal, it does so as a demonstration of excellence rather than conformity to local type. The beautiful (mostly late) symmetrical handaxes were never bog-standardly typical of their place and time. These considerations apply to signals within a community. Perhaps instead conformity in craft technique functioned to send signals between communities. Uniformity in craft technique sends a message of social cohesion to potential allies and enemies, creating an interest in conforming, and in others' conforming. It is true that if the intercommunity environment is fraught, it can be in the community's interest to send a convincing signal of coherence and unity, advertising the fact that they are good allies and dangerous enemies. Expensive rituals are often explained as expensive hence honest signals of community cohesion (Sterelny 2020b). However, as Steven Kuhn and Mary Stiner point out, signals between communities must be easily understood and of high amplitude (Kuhn/Stiner 2007). For they must be read at a distance. A foraging party all carrying identical oblong shields with a large red eye in the centre satisfies these conditions. Carrying a handaxe knapped from a core struck from the left of a bedding plane does not.

So here is an alternative suggestion as to how the scope of intrinsically motivated, standard-guided actions might have extended to social interaction, not just to individuals regulating their own craftwork. I agree with Jonathan's idea that shared activities are crucial. So natural bridges are shared foraging activities, like hunting. Hunting is a highly skilled activity. Ambush hunting, likely to have been an early and important form of human hunting, is collaborative, requiring team work. It is intolerant of error. Those positioned in wait have to be quiet and still, often despite serious physical discomfort (Pickering 2013). These discomforts are vividly described in Richard Gould's description of a failed emu hunt in the Western Desert: the wait was long; the sun was hot; the flies and ants were many and voracious (Gould 1969). These forms of collaborative action have elements both of craft skill and social interaction. There is a way of getting a hunt right, which has both social elements—smooth coordination, seamless division of the spoils—and precise physical execution. Moreover, this pathway to other-directed normative regulation explains the intensity of affective emotions directed at others' action. It explains the strength of approbation or disapproval directed at second and third players. Collective hunts that miscarry because of a stuff-up really hurt. They impose costs on all participants. Likewise, those that succeed because of exceptional skill deliver a real benefit. In contrast, it is hard to explain why an agent should care profoundly about others' unusual techniques for thinning a blank.

In short, I think it is quite likely that the normative regulation of action began with craft skills, and with agents developing conceptions—standards—of well-made tools and of skilled toolmaking. It is likely that norm-guided action began with the evolution of pride in smooth fluent performance, and discontent with clumsy execution and clunky tools. My guess that this extended to admiration of other's adept performance, and annoyance at inept performance, in the context of those collective actions which required competence and care from all the participants. These are contexts in which an agent has reason to care about the performances of others. Moreover, these emotions, positive and negative, would be apparent both to their targets and to bystanders, and the importance of reputation gave these audiences a reason to notice and respond. So these other-directed emotions were likely to have social consequences. This affective response to others' actions is driven by their material consequences to you. In collective action, others' skills (or lack thereof) shape your rewards. On this picture, norms of skill were likely to have played an important role in the Acheulian and in the Middle Stone Age, as yet more demanding stone crafting began to establish about 500 kya. These norms became other directed in the specific context of skilled collective action. In the middle Pleistocene, the scope of these other-directed norms was quite limited; to one aspect of the subsistence economy, albeit a central one. Other-directed, impersonal, more overtly social norms became important in the more socially and economically complex worlds of the last 100 k or so years of the Pleistocene (Sterelny 2014; 2020a; 2021).

In short, yes to the skill hypothesis, but in a form which emphasises the role of proto-normative affect in the diachronic development and deployment of skill. Moreover, I see these normative emotions as becoming other-directed in the context of skilled collaborative foraging, rather than in individual but public tool making, with each becoming aware of, and caring about, others' craft habits.

References

Birch, J. (2021), Toolmaking and the Origin of Normative Cognition, in: *Biology & Philosophy* 36, 4

Boyette, A. (2016), Children's Play and Culture Learning in an Egalitarian Foraging Society, in: *Child development* 87(3), 759-769

Boyette, A./B. Hewlett (2018), Teaching in Hunter-Gatherers, in: *Review of Philosophy and Psychology* 9, 771-797

Fodor, J. A. (1983), *The Modularity of Mind,* Cambridge/MA

Gould, R. A. (1969), *Yiwara: Foragers of the Australian Desert,* Sydney-London

— (1980), *Living Archaeology,* Cambridge

Hewlett, B./H. Fouts et al. (2011), Social Learning among Congo Basin Hunter–Gatherers, in: *Philosophical Transactions of the Royal Society of London B: Biological Sciences* 366(1567), 1168-1178

Hiscock, P. (2004), Slippery and Billy: Intention, Selection and Equifinality in Lithic Artefacts, in: CAJ 14(1), 71-77

— (2014), Learning in Lithic Landscapes: A Reconsideration of the Hominid 'Toolmaking' Niche, in: *Biological Theory* 9(1), 27-41

Kuhn, S. (2020), *The Evolution of Paleolithic Technologies,* London

— /M. C. Stiner (2007), Body Ornamentation as Information Technology: Towards An Understanding of the Significance of Early Beads, in: *Rethinking The Human Revolution,* P. Mellars/K. Boyle/O. Bar-Yosef/C. Stringer (eds.), Cambridge, McDonald Institute for Archaeological Research, 45-54

Lew-Levy, S./N. Lavi et al. (2018), How Do Hunter-Gatherer Children Learn Social and Gender Norms? A Meta- Ethnographic Review, in: *Cross-Cultural Research* 52(2), 213-255

— /R. Reckin et al. (2017), How Do Hunter-Gatherer Children Learn Subsistence Skills? A Meta-Ethnographic Review, in: *Human Nature* 28, 367-394

Naveh, D. (2016), Social and Epistemological Dimensions of Learning Among Nayaka Hunter-Gatherers, in: H. Terashima/B. Hewlett (eds.), *Social Learning and Innovation in Contemporary Hunter-Gatherers: Evolutionary and Ethnographic Perspectives*, Dordrecht, 125-134

Pickering, T. R. (2013), *Rough and Tumble: Aggression, Hunting and Human Evolution*, Berkeley

Shipton, C. (2019), The Evolution of Social Transmission in the Acheulean, in: K. Overmann/F. Coolidge (eds.), *Squeezing Minds from Stones: Cognitive Archaeology and the Evolution of the Human Mind*, Oxford

— /M. Nielson (2015), Before Cumulative Culture: The Evolutionary Origins of Overimitation and Shared Intentionality, in: *Human Nature* 26, 331-345

Sterelny, K. (2014), A Paleolithic Reciprocation Crisis: Symbols, Signals, and Norms, in: *Biological Theory* 91(1), 65-77

— (2019), The Origins of Multi-Level Society, in: Topoi 40, 207-220

— (2020a), Norms and Their Evolution, in: *Handbook of Cognitive Archaeology*. T. Heneley/M. Rossano/E. Kardas. (eds.), New York, 375-397

— (2020b), Religion: Costs, Signals, and the Neolithic Transition, in: *Religion, Brain & Behavior* 10(3), 303-320

— (2021), *The Pleistocene Social Contract: Culture and Cooperation in the Hominin Lineage*, New York

Stout, D./M. Rogers et al. (2019), Archaeology and the Origins of Human Cumulative Culture: A Case Study from the Earliest Oldowan at Gona, Ethiopia, in: *Current Anthropology* 60(3), 309-340

Terashima, H./B. Hewlett (eds.) (2016), *Social Learning and Innovation in Contemporary Hunter-Gatherers*, Dordrecht

Wilkins, J. (2020), Learner-driven Innovation in the Stone Tool Technology of Early Homo sapiens, in: *Evolutionary Human Sciences* 2, e40

Wrangham, R. (2017), Control of Fire in the Paleolithic: Evaluating the Cooking Hypothesis, in: *Current Anthropology* 58(S16), S303-S313