Peter Railton*

Normative Guidance, Evaluative Guidance, and Skill

https://doi.org/10.1515/auk-2021-0014

Abstract: At least since Aristotle, practical skill has been thought to be a possible model for individual ethical development and action. Jonathan Birch's ambitious proposal is that practical skill and tool-use might also have played a central role in the historical emergence and evolution of our very capacity for normative guidance. Birch argues that human acquisition of motor skill, for example in making and using tools, involves formation of an internal standard of correct performance, which serves as a basis for normative guidance in skilled thought and action, and in the social transfer of skills. I suggest that *evaluative* modeling, guidance, and learning play a more basic role in motor skill than standards of correctness as such—indeed, such standards can provide effective normative guidance thanks to being embedded within evaluative modeling and guidance. This picture better fits the evidence Birch cites of the flexibility, adaptability, and creativity of skills, and can support a generalized version of Birch's 'skill hypothesis'.

Keywords: normative guidance, evaluative modeling, learning, ethics, skill, tool-use, evolution

1 Introduction

As a philosopher with no special expertise on early hominid evolution, what I find especially exciting about Jonathan Birch's ambitious paper is that he seeks to make more concrete the analogy often made in the history of philosophy between ethics, on the one hand, and skill or craft, on the other. Of course, the analogy has typically been thought to be partial—the imperatives of skill or craft have been seen as 'technical', 'instrumental', 'hypothetical', or 'contingent', while those of ethics have been seen as 'rational', 'non-instrumental', 'categorical', or 'obligatory'. But the analogy is nonetheless held to provide a potential model for some key aspects of the cognitive and practical capacities involved in everyday ethical life,

***Corresponding author: Peter Railton,** Department of Philosophy, University of Michigan, Ann Arbor, USA, email: prailton@umich.edu

including: the need for individual experience and social learning, the possibility of open-ended understanding and competence difficult to capture in determinate principles or rules, the existence of non-deliberative or intuitive action-guidance in ways that are nonetheless sensitive to a complex array of considerations, and so on. If we could better understand the perceptual, cognitive, and practical capacities at work in skill, this might help remove some of the mystery about the perceptual, cognitive, and practical capacities seemingly at work in ethics.

Birch makes an important contribution to developing such an understanding by bringing into the philosophical discussion a more empirically-grounded account of skill combined with a more clearly-formulated set of conditions for normative guidance. And he thereby is able to push the analogy into new territory, by suggesting how skill acquisition, deployment, and transfer might be *generative* of capacities for normative guidance. Rather than confining himself to the claim that normative guidance can be *likened* to skill, he argues that complex motor skills, and the tool-making and tool-using practices and traditions they help make possible, by their nature *involve* normative as well as technical capacities, and may be a key component of the story of the evolution of humanity's distinctive forms of life.

This strikes me as a highly promising idea. There is, once one has given a more plausible account of skill and normative guidance, so much about the development and teaching of motor skill that could help underlie and fund the emergence and development of human normative thought and practice that it seems likely that motor skill will indeed be an important part of the story of how we became beings whose existence is, as Birch puts it, 'guided by thousands of norms'.

I will begin by considering Birch's account of skill, and then his account of normative guidance. After that, I will pursue a few related questions about these accounts and how they fit together. My comment will thus be limited to the very first steps in Birch's account. But my hope is to contribute to the further development of his promising idea.

2 Skill and Skill Acquisition and Deployment

Philosophical thinking about skill has been heavily influenced by psychological accounts that emphasize what appear to be the 'automatic' or 'subpersonal' perception-action links acquired when someone acquires a skill. A complex activity that initially is cognitively and practically effortful to perform, requiring on-going conscious attention and deliberative guidance, can through training and experience become something that can be accomplished with minimal self-conscious higher-order cognition or deliberation, freeing skilled agents to attend to other things.

This view of skill surely owes its popularity in part to its fit with the phenomenology of skilled action—a teacher who is also a skilled bicyclist, and who is running late for class, can spend virtually the entire rushed ride from home to school mentally composing her lecture, with no need to concentrate on staying upright or steering through curves. Birch observes that some have even characterized the exercise of skill as relatively 'mindless'—the agent need not have anything in particular in mind in order to effectively deploy a well-developed skill. However, it is difficult on this account to explain the 'on the fly' flexibility and creativity of highly-developed skills, as seen in the evident capacity of top athletes, musicians, or players of strategic games to improvise highly effective novel behaviors under the pressure of time. Even if we imagine a skill to be constituted by a wide array of basic subroutines, each of which is 'automatic' but which can be reassembled in novel ways to generate flexibility and adaptability, still, there would need to be a higher-order mental component of skill able to innovatively and intelligently call upon and sequence these subroutines 'on the fly', with little or no explicit self-conscious deliberation.

In order to account for the flexibility of skill, psychologists and cognitive scientists have argued that skilled individuals acquire implicit mental models of situations and actions, capable of simulating actions and outcomes in real time, generating predictions, and guiding behavior accordingly. And to account for the intelligence of such non-deliberative guidance, they posit feedback mechanisms through which these models are spontaneously updated in response to gaps between predicted and actual outcomes.

Birch adopts a particular version of this sort of model-based motor control, designed to help explain how higher-order agential thought can enter into the operation of well-trained skills. In addition to the sub-personal modeling, simulating, and motor planning that takes place in cerebellum, which Birch sees as 'largely encapsulated from cognition', there is also a cortically-based 'cognitive control model' that is at least partially consciously accessible, so that agents can introduce novel or abstract information, situational awareness, and revised strategies 'on the fly'. Drawing upon the work of Wayne Christensen and colleagues (Christensen et al. 2015; 2016), who studied such complex motor skills as competitive mountain-biking, Birch gives the following characterization of the nature and role of the higher-order 'cognitive control model':

[1] 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" (193; [quoting] Christensen et al. 2015, 344).

He further explains:

[2] 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. [1]

Because these models are predictive, they support feedback learning and can become increasingly 'well-calibrated' through experience.[2]

Using the example of mountain-biking, Birch writes:

[3] 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). (Birch 2021, 193)

Those with more highly developed motor skills have more highly developed and well-calibrated mental models, but because the higher-order model is not merely sub-personal and cognitively 'encapsulated', updating can occur through more conscious means, and this allows a role especially for *social learning*.

We now begin to see an important connection between culture and skill, since expertise painfully acquired by others over the course of *their* lives can be transmitted through communication, instruction, or observation to enrich one's own cognitive control models—and thereby one's skill level. This promises to be an evolutionary 'virtuous cycle': social learning can help generalize a useful skill in a population, thereby bringing yet wider experience and greater cognitive power

1 Although Birch mentions sensory feedback, most fine motor control in real time must take place too rapidly for sensory feedback to be the primary 'learning signal'. Instead, it is thought that the brain directly compares an efferent copy of the motor commands it issues with its forward model, and uses discrepancies generated in this comparison to update the model and subsequent motor commands

2 Birch at one point considers the possibility that a cognitive control model might have a larger role in the early stages of skill acquisition, rather than in skill's most fully-developed stages.

to skill development, which in turn could improve levels of food production, provision of shelter, group defense, etc., which then could sustain higher rates of reproduction, which in turn could lead to further expansion and development of skill, including the emergence of a division of labor and greater specialization, and so on. However, for this social process to operate there must be an individual psychological infrastructure that subserves effective learning and teaching of skills via these social means. If psychological variations occurred that enhanced such capacities or motivation for teaching and learning, these might be selected for and play a significant role in the rapid progression of human evolution over the short span of *Homo sapiens*' existence. Skilled tool-making and tool-use, with their capacity to dramatically increase productivity, could be a central locus in the building of human culture.

Anthropologists sometimes call modern humans 'obligate tool users' in contrast with our primate relatives and hominid ancestors. Through what appears to be a history of *bio-cultural evolution* involving tool-use (including such 'tools' as controlled fire), human *physiology* has ended up with distinctively smaller teeth and jaws, shorter digestive tracts, more limited physical strength and climbing ability, little natural protection from the elements, a brain with very high developmental and metabolic energy costs, and an extended period of dependent infancy. Humans would be lucky to survive and reproduce at all if stripped of the totality of our tools. Might the acquisition and transfer of tool-making and tool-using skills have had a similarly profound effect on the bio-cultural evolution of our distinctive human *psychology* and *culture*? What kind of psychological and cultural infrastructure might favor the acquisition and transfer of complex motor skills? Birch's claim is that it would involve not only capacities for, and practices of, imitation or sociability, say, but also for 'normative cognition' or 'normative action-guidance'. This is the second half of his 'skill hypothesis':

> [4b] Some 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." (192)

3 Normative Cognition and Normative Guidance

Philosophical thinking about normative guidance of cognition and action has been largely dominated by accounts that focus on agential embrace of, identification with, or commitment to norms—it is this kind of agential involvement that explains how we come to be guided *normatively*. That is, what makes guidance by a rule *normative* for the individual—and not a matter of mere habit, imitation, conformity, or social enforcement—is that compliance with the rule is something flowing *from* the agent and her own volition, not something that simply happens to her or is imposed upon her. It might be thought that, as infants, we initially conform to certain rules of behavior because this pleases adults or escapes their anger, but only once we can begin to make up our own mind about the rules, and impose them upon ourselves without need for external enforcement or reinforcement, does fully-developed normative guidance emerge. Like the relatively 'mindless' conception of skill, this highly 'mindful' conception of normative guidance has a seeming phenomenological plausibility. At the same time, however, the seeming mismatch between more developed skill being increasingly mindless, and more developed normative guidance being increasingly mindful, would seem to make skill an implausible choice to capture what it is distinctive about normative guidance in ethics after all.

However, once again, what appears to be a phenomenologically-plausible account runs into problems explaining actual practice. First, we know from observing such phenomena as language and conversation that individuals can be normatively guided by large constellations of norms that they could not articulate, much be understood as endorsing or applying deliberatively. However, these norms are *not* for us mere habits, or externally-enforced rule-following. Rather, they encode our sense of how one *ought* to speak, and we regard departures as *mistakes* or *faults*, not merely anomalous or unwanted events. Moreover, these rules are not merely imposed upon the individual through social sanctions—in many societies, very little explicit language instruction and correction occurs in the first years of life, yet infants become entirely fluent. Such implicit norms might be put aside by philosophers as something less than 'full-blooded normative guidance', in contrast to norms explicitly adopted through agential endorsement and commitment, etc., and thus not of interest for ethics. Yet much of ordinary morality appears to have a similarly implicit character—and to emerge in child development without explicit instruction and sanctions. While controversy persists over timing, evidence suggests that, by 18 months, infants already spontaneously prefer third-party helpers over third-party hinderers, and actively take steps to help unrelated others and to share the results of joint activity fairly without any external instruction or reinforcement (Warneken/Tomasello 2006; Schmidt/Sommerville 2011). By age 3-4, infants already spontaneously distinguish between situations in which an adult in authority imposes an arbitrary but harmless rule, which they may willingly obey, and situations in which an adult in authority is imposing a harmful 'rule', which they often willingly resist obeying, again, without any adult instruction or reinforcement and even in the face of unwanted disapproval from the adult in authority. Neither is such refusal to obey 'mindless' habit or dislike—indeed, if queried, infants will explain why they won't comply in terms of preventing harm or unfairness (Turiel 2008; Smetana et al. 2012). Such behavior, which can be observed in infants in a range of societies, strongly suggests emerging competence in normative guidance both for selves and for others, which seems to parallel the development of theory of mind (Wellman 2014). The guidance is not simply coming 'from without' and without understanding of its point or value, but 'from within', and linked with some understanding of its point or value. In fact, this looks quite a bit like a flexible, situation-sensitive emerging *skill* with the ethically-relevant dimensions of situations and actions, akin to the infant's emerging flexible, situation-sensitive skill with language and communication.

How might such forms of intelligent normative guidance be characterized, according to Birch? Here are the core components of his answer:

[An agent] could never verbally express the full set of norms that guides [their] behaviour. In such cases, three key psychological ingredients are in place …

[i.] First, the agent reliably *notices* or *anticipates* failures to comply with the norm, in themselves or in others.

[ii.] 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.

[iii.] 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.

As I casually converse with others, I am not aware of—and even with reflection would be unable to articulate—all of the linguistic, interpersonal, cultural, epistemic, and ethical norms I am following. But someone observing my behavior could see them at work in the 'affective pressure' I feel when they are violated. I might, for example, evince a clear uneasiness when a conversation partner comes closer than a certain 'normal conversational distance', and even take a step backwards to 'restore conformity' with the conversational norm, all without consciously attending to this. Someone from a different culture, having internalized different norms of conversational distance, would behave quite differently in the same situation.

4 Connecting Skill and Normative Guidance

We now can make more definite the connection between complex motor skills and normative guidance. We can see the elements [i]-[iii] at work in skilled behavior because, according to Birch, the cognitive control model that participates in the guidance of skilled behavior 'implicitly encodes a standard of correct action', that leads the agent to notice (if only implicitly) departures from the norm, to feel corresponding 'affective pressure' to 'prevent or correct the departure', and 'know what to do' to 'restore conformity'.

How do we identify this 'standard of correct action'? As we saw in the [i]-[iii], the working of normative guidance via a standard will be most readily seen in cases of 'mismatch' between standard and behavior. Moreover,

> [5] … 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. … 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.

Birch further explains the normative character of this implicit 'standard of performance':

> [6] Skill leads to discontent when the agent falls short of the standard of performance implicitly encoded in the control model. An incorrect adjustment, leading to a mismatch between the predictions of the cognitive control model and the agent's behaviour, *feels wrong* to the agent, independently of (and often temporally prior to) any physical discomfort the error may cause. Skill creates internal pressure to conform to an internalized standard of correct performance.

Thus, for Birch, complex skill has an inherently norm-guided aspect. If not already implicit in [i]-[iii], we can add, within the spirit of Birch's idea:[3]

iv. Fourth, the affective pressure arising from mismatches is not merely attributable to unwanted or disappointing further consequences of the mis-

3 He mentions, for example, a connection between aptitude for complex skills and *intrinsic* motivation for mastery.

matching behavior, rather, it arises in part *from the mismatch as such*, without any further incentive.

Under conditions [i]-[iv], we can see that a pattern is functioning as a *standard* for the agent—so that a mismatch is experienced as *in itself* an error by her own lights. The mismatch itself '*feels wrong*' and constitutes an 'error' in the eyes of the agent, even when it brings no further negative effects. Birch writes:

[7] 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.

This now gives us the first half of Birch's 'skill hypothesis', to be combined with [4b]:

[4a] 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.

So if evolution favored a mind capable skilled motor control, it also favored a 'basic platform' for guidance by norms. As we noted, the norms may be 'technical'. But the motivation involved is not merely 'instrumental'—to some extent there is intrinsic motivation to follow the norm. Some primatologists have argued, for example, that chimpanzee imitation is always limited to the 'instrumental'—once the chimp sees the goal achieved, it is said to be uninterested in any aspect of the process leading up to the goal that is not directly instrumental. Human infants, by contrast, will imitate the process even in its non-instrumental dimensions (Carpenter/Call 2009). They are 'modeling their action' on the process, treating departures from the process as 'to be avoided' as such—not for the sake of something else. And this is why a normative capacity meeting conditions [i]-[iv], even if originally the *content* of the norms acquired were 'technical', would constitute a *competence in normative cognition and guidance* that could be a foundation for guidance by norms of a less technical kind. Thus Birch can write:

[8] … 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. (199)

5 Questions

If this sketch of the first steps in the development of Birch's view is roughly right, I'd like to pursue somewhat further two questions about the role of norms in motor skill as he understands it, and how this might relate to questions about the infrastructure and origins of human normative psychology more generally.

My first question concerns Birch's claim that the cognitive control model involved in complex motor skill "encodes an implicit standard of correct action" that serves as the basis for the normative guidance involved in skill. Passages [1], [2], and [3] make it clear that a cognitive control model is *descriptive* in character. It is "a representation of the causal structure of a complex skill and the situation in which it is executed", akin to a "causal graph". As such, it would have *mind-to-world* direction of fit—the relevant notion of 'incorrectness" for such a model would be *representational error* or *falsehood*. This might be manifest, for example, when there is a 'mismatch' between what the model predicts as the causal consequence of a given action in the present circumstance and what actually happens when the agent performs that action in this circumstance.

A standard of correct behavior, by contrast, does not purport to describe or predict behavior, and so is not 'incorrect' or 'false' if actual behavior does not conform to it. Rather, standards are said to have *world-to-mind* direction of fit in that they set a condition that behavior *is to meet*. Other things equal, when an action does not fit the norm, the *normative error* or *fault* is attributed to the action, not the norm.

Now Birch doesn't claim that a norm of correct performance is *explicitly* represented in the cognitive control model. But, given his description of the cognitive control model, I do not see how such a standard could be *implicitly* represented in or implied by the model. A model of actual or possible causal relations for actions and circumstances would appear to be silent on the question of which actions are the ones *to be chosen* by the agent.[4]

Let's consider the example of competitive mountain biking used by Birch. Birch writes, quoting studies by Christensen et al. (2015; 2016), that the cognitive control model:

---

4 Some mental states—like *affective* or *evaluative* states—can have both directions of fit. For example, an affective state like fear simultaneously presents a situation in a way that could be more or less accurate with respect to the magnitude or sources of risk, and motivates and helps guide a suite of risk-relevant responses to it (Railton 2017). But this does not seem to be true of causal representations as such.

[9] … mediates between explicit plans and low-level (cerebellar) motor control, by 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". (193)

Once the picture of the operation of motor control has been filled out in this way, we *can* see something with world-to-mind direction of fit—independently-specified 'performance goals' of the rider have been included. Relative to these goals, the rider can fail or succeed by her own lights. But this kind of 'discrepancy' in performance is an example of the familiar means-end normativity of goals for actions, and does not show that a 'standard of correct performance' lies within the cognitive control model itself.

Another way to see this is to consider that the biker might, on a given run, have 'performance goals' different from smooth riding, and yet still draw upon exactly the same causal model "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)" [3]. Perhaps the biker has been stung by criticism that she's 'too careful' to be an exciting biker, despite her victories, and so she decides on this run to cycle 'right on the edge', achieving the wildest ride compatible with staying in one piece—to show that she is skilled enough to do this. Now the fact that she comes into a curve too fast to round it smoothly, and instead skids impressively and regains control only at the last instant, would enable her to *satisfy* her performance goals, while smooth rounding around the curve would discomfort her as a performance error—'Rats. I'll have to speed up into the next curve.' So neither performance goal, smooth riding or wild riding, is implicit in, or contrary to, the cognitive control model.

To be sure, the cognitive control model is said to "represen[t] those aspects of the causal structure of the situation relevant to successful execution of the skill, anticipat[ing] upcoming obstacles and problems, and predict[ing] the flow of sensory feedback that will occur if skill execution is successful" [2]. This can make it look as if some criterion of 'successful execution' is part of the causal model as such. But 'successful' here is a variable term—skill is impressive in part because, for a broad range of possible goals and circumstances, the skilled individual can successfully pursue *whichever* of these goals they have. If the skill is in the use of a tool, like a mountain bike, then highly skilled individuals will have a wide repertoire of goals they can use the bike to accomplish: riding smoothly, riding wildly, riding in adverse conditions, riding under great uncertainty, riding to impress a selection committee, or riding to finish with a time just behind a friend who is recovering from a serious fall and whose spirit they'd like to restore—while *appearing* as if trying to win.

This brings us to a general feature of motor control—a motor controller won't be well-defined without something like a target value, goal specification, cost-function, or reward-function. So what lies within a 'cognitive control model' capable of governing motor behavior would be not only a causal model of the system to be controlled (myself, the workpiece and tool, etc.), but also a goal or value function of some kind. Not necessarily a fixed goal or value function, but a capacity to assign value such that, when a goal is specified, this can be translated into a *gradient* that effectively guides motor responses in real time. After all, fit with a norm provides two kinds of information—fit or failure to fit. But skilled motor control needs to be continuously guided to *approach* its target without over-compensating or falling short, without running out of energy, etc. Consider a master flake-stone tool-maker setting before the learner a sample hand axe. Stones vary indefinitely in their composition and structure, so the learner needs a control model that can take the specified task of reproducing this sample and use it to *regulate* the ways she strikes the workpiece, assigning a positive value to approaching the shape of the sample and a negative value to time spent or materials consumed. The sample has become something like an *ideal* to be approximated as well as possible. But even with infinite care, differences inherent in the stone will mean that the ideal can only be more or less well-approximated. The learner's value function nonetheless needs to assign value such that there is continuing guidance and motivational encouragement along the way. Failure to precisely fit a norm is simply that, failure. Individuals seeking to attain a goal, by contrast, should not read an action leading to a discrepancy with the internal predictive model that constitutes a *better*-than-expected result as nothing more than an error to be corrected—rather, such events should serve to update value assignments to actions within the model.

In general, effective motor control in natural and artificial systems is inherently linked to *learning*, and there must be reward for movements in the direction of the goal even when one is far from it, and so not in accord with a 'standard of correctness'. Systems theory and a variety of evidence from highly-skilled motor behavior suggests that avoiding 'overlearning' or 'automatization,' which can make motor behavior too rigid to track changing situations or newly-emerging opportunities, gives motor variability and exploration a role to play even in expert performance (Davids et al. 2003; Bartlett et al. 2007). Neuroimaging evidence suggests that elite athletes show high levels of brain connectivity in which key components of motor control can recruit information widely, and integrate multi-sensory signals and explicit information in real time (Huang et al. 2017; Krakauer/Mazzoni 2011; Van Overwalle et al. 2014; Hull 2019). While this might count against the idea that cerebellar motor planning and control are 'cognitively encapsulated', it fits even better with Birch's larger picture of skill as involving elements of 'cognitive control' rather than mere 'automatization'.

Typically in acting we have multiple goals or multiple values at stake, and there will be multiple ways of realizing these to greater or lesser degrees. Going back to image of casually conversing with another, I might be seeking to communicate, to learn, to warn, to solidify a relationship, and to do as I ought, all at once. Moreover, I must do this while having no direct control over the others in the conversation, and so must be prepared to engage, accommodate, adjust, encourage, and so on, if I am to be successful to a reasonable degree. The ability to be *resilient* and *adaptive* in these ways is part of what we think is involved in social and linguistic skills—their guidance of action is not 'brittle', and should allow recovery from missteps and learning-along-the-way. Birch writes, of the normative guidance provided by skill:

[10] Skill execution must take a very specific course (the skill must be executed 'just the right way') to avoid triggering any dissatisfaction.

However, even if we take a very tightly-constrained instance of complex motor control, say, mountain bikers with the overriding goal of descending the mountain in the quickest time, two bikers of consummate skill could end up tied for the lead despite many differences in the precise adjustments they made while descending to pressure on the pedals or brakes, angle of steering, distribution of weight, etc. As Birch's discussion of Christensen et al. (2015; 2016) indicates, competitive bikers often are allowed little time to study the trail or the bicycles they will be using—so their planning is perforce partial at best. Skill will help them fill in the plan along the way, but there won't typically be a 'just the right way' to do so. The way that skill regulates behavior thus should incorporate capacities for guidance of behavior in light of multiple values or ends at stake, and making needed trade-offs, accommodations, compensations, innovations, explorations, etc., in real time, rather than giving us an *a priori* standard of 'just the right way' of acting.

My second question concerns whether the idea that complex motor skills involve at base this kind of causal-evaluative modeling and control, rather than an internalized 'standard of correct performance', would do anything to undermine Birch's larger project of looking to skill acquisition and use as a potential base upon which further normative guidance could be built. My sense is that it would not—and perhaps, on the contrary, it would make it clear that the 'skill hypothesis' is more robust with respect to particular mechanisms that might instantiate it. Consider conditions [i]-[iv] again, put in terms of evaluative guidance:

i'. The agent reliably *notices* or *anticipates* departures from an evaluative goal or ideal.

ii’. The agent feels *affective pressure* (for example, in the form of discomfort, shame or anger) to approach toward the goal or ideal, or prevent movement away from it.

iii’. The agent *knows what to do* to make progress toward the goal or ideal in a way appropriate to the situation. This may involve changing their own behaviour, trying to change the behavior of another, seeking forgiveness, or administering punishment.

iv’. The affective pressure arising from failure to make progress, or to match expectations of progress, toward the goal is not merely attributable to unwanted or disappointing further consequences of the lack of progress or mismatching of expectations, rather, it arises in part *from the lack of progress or mismatch with expectations as such*, without any further incentive.

How might *norms* or *standards* come into this picture and become part of our distinctive forms of normative competence? Norms and standards cannot in themselves generate motivation or behavior. They can play the role they do in the guidance of human behavior—can ‘come to life’ in how we think and act—because creatures like ourselves can implicitly or explicitly take an internally represented norm or standard as a goal or ideal. By approaching more closely to norms or standards we have internalized, we can become more predictable and consistent over time as individuals and across individuals, and evaluative modeling can supply the guidance and motivation for this—even when we are far from the norm or standard. This motivated capacity to approximate norms or standards in turn facilitates activities involving interpersonal coordination, collaboration, or communication, by furnishing something like a ‘common ground’. For example, if there is a task that requires multi-agent collaboration, then norms that specify particular roles can enable us to work together with reasonable expectations of the others’ behavior and a clearer view of what they expect from us in return. These roles can guide motor behavior by serving as ideals or ideal-types independent of the individual filling the role, and each be associated with a particular value function to guide and reward approximation to the role ideal in practice (compare Tomasello 2016). Given [i]-[iv], above, there will be *internal motivation* to fulfill the role, or to hold others to the role, e.g., internal affective pressure against failing to do one’s part, even in the face of other, opportunistic incentives toward free-riding or defection. From this can arise ideas of what it is to act well in a joint endeavor, and how it is a valuable characteristic—a virtue of ‘fairness’ or ‘trustworthiness’—to be reliably motivated in this way to do one’s share in joint project and to help others to do their share. At the same time, this internal affective pressure disposes one *symmetrically* to disapprove and correct oneself as well as others when this fails to occur. Norms or standards thus can be taken up *as normative* in an evaluatively-guided normative psychology, and come to function explicitly or implicitly as *shared ideals* or *shared standards* against which we can do better or worse, individually or together, and care about how well we do.[5]

Thus, once such a cooperative or collaborative psychology is in place, the *content* of the ideals, and the associated value functions, can vary while the control processes remain the same. [i]-[iv] can function as a 'basic platform' that can serve not only technical norms, but "more abstract norms of reciprocity, ritual, kinship and fairness" [8]. Generalizable skill at identifying and following norms seems to appear early in human infants (Aslin/Newport 2012), and they can quickly get the hang of rule-based games. Thanks to such generalizable skill, we can understand how it is possible for humans to be able to negotiate the elaborate structure of norms that help guide our lives and make our distinctive, large-scale yet cooperative social existence possible. And thanks to underlying human value structures, we can understand how systems of cooperative or collaborative norms can motivate humans and engage their participation without external enforcement.

One of the longest-running role-based cooperative or collaborative human practices processes guided by shared or overlapping aims or values, and motivated independently of external enforcement, is surely teaching and learning. In its more developed forms, this practice involves the teacher and learner recognizing the situation, knowledge, and needs of the other, and developing a form of joint agency involving mutual accommodation to permit effective communication of what is to be taught. One suspects that evaluatively-guided teaching and learning practices of this kind emerged in humans for an array of reasons, given how much human infants must learn in order simply to manage in the world. However, tool-making and tool-use would be a natural locus for such teaching and learning practices to develop. Indeed, transfer of tools from teachers to learners, along with instruction in tool-making and tool-use, can constitute a direct transfer of a 'package' of accumulated experience or culture that might not need words, the transfer of which can be observationally monitored, and that can increase productivity more than it costs. For such incentives to work most fully and reliably, problems of the distribution of increased productivity need to be solved, as Birch suggests. Complex teaching and learning take substantial time away from hunting and

5 Birch develops an example akin to this in the case of standardization within a division of labor in hand-axe making. My thought would be that the skill underlying this capacity for standardization is itself evaluatively guided—that the standard is 'brought to the skill' by the particular purposes of the individuals participating in this division of labor. Such standardized 'template matching', given the non-standardized basic ingredients, is indeed a skill—a capacity made possible by a 'basic platform' of flexible evaluative guidance that can take a template as a shared standard to be approximated.

foraging (Pargetter et al. 2019), and if this loss is not to be concentrated precisely in the individuals showing the increased normative competence involved, there will need to be background social practices able to support this—such as the mandatory food-sharing norms observed in many contemporary hunter-gatherer societies (Boehm 2012).[6]

---

6 Tool-making and tool-use, and conveying such tool-making and -use across generations, would not seem on their own to bring such a background into existence, since various apes engage in these practices without developing wider practices of normatively guided sharing, or joint intentionality in teaching and learning. Yet an intriguing recent study (Musgrave et al. 2021) of chimpanzee tool use and teaching in the wild provides some fuel for imaginative speculation about how the 'skill hypothesis' might help with the 'chicken-and-egg' questions Birch considers. In both the Goualougo Triangle in the Republic of Congo and in Gombe in Tanzania, chimps have developed tools for fishing termites out of mounds. In Goualougo, tool use is more complex, involving multiple types of tools, the selection of specific materials for different tools, sequential utilization of the tools, and so on; in Gombe, only one type of tool is used, made of no special material. – What Musgrave and colleagues observed in this small comparative sample is that, in the chimpanzee culture in Goualougo with more complex tool use, there was also a higher rate of cooperative behavior, with much greater willingness of chimp mothers to transfer tools to their infants when requested. Such non-verbal transfers of a 'package' of accumulated culture constitutes a form of teaching, and comes at some cost to the mother, who loses the 'congealed labor' in the tool. This transfer thus isn't explained in terms of conspicuous immediate instrumental gains (Musgrave et al. 2021), although the mother presumably has a robust evolutionary interest in promoting the ability of her offspring to provide adequate nutrition for themselves, so this teaching practice would not simply be penalized evolutionarily. It seems, therefore, that a kind of longer-term, non-reciprocated helping and teaching behaviors—and individual value-functions that could guide them—have gotten some foothold in Goualougo, enabling cross-generational continuity of a culture of more complex tool-making and tool-use. Over time, perhaps the example of mothers' willing responses to requests for tool transfer might encourage in infants a sense that this is what is expected, without any conspicuous instrumentality, and so something like a role-related proto-norm of teaching could get underway and carry on over time, partially subsidized by the greater productivity made possible by complex tools. Complex tool use, and the teaching and learning practices that can preserve such use over time, might thus be a locus for the development of a proto-normative culture the productive advantages of which could in principle favor selection for a psychology that more readily acquires such culture—akin to the 'virtuous circle' described earlier. his is wandering much further into the thickets of evolutionary theory than I have any right to, but it suggests that a form of Birch's 'skill hypothesis' might be part of the story of the evolution of capacities that could help initiate movement in the direction of a more genuinely normative psychology even clearly before the emergence of language and modern humans. After all, as Birch notes, language understood as a norm-guided, mutually-accommodating enterprise seems to *presuppose* something like a proto-normative psychology, while the 'skill hypothesis', starting from simpler, less intentional, and less expressly normative ingredients, could help explain the emergence of such a psychology. Progress in tool-making and tool-use, and development as normatively-guided creatures capable of working together to create and communicate much more

So, even though I have raised questions about whether we should see a 'standard of correct performance' as internal to skill as such, I believe we can arrive at an (only slightly) edited version of Birch's 'skill hypothesis':

> 4a'. In modern humans, complex motor skills and craft skills, such as skills related to toolmaking and tool use, are guided by *forms of evaluative control that can afford a 'basic platform' for the development of normative guidance more generally*.
> 4b'. Some core components of human normative cognition *therefore may have* 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.

Vindicating what he calls his 'guiding thought':

[11] … 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.

References

Aslin, R.N./E. L. Newport (2012), Statistical Learning: From Acquiring Specific Items to Forming General Rules, in: *Current Directions in Psychological Science* 21, 170-176

Boehm, C. (2012), *Moral Origins: The Evolution of Virtue, Altruism, and Shame*, New York

Carpenter, M./J. Call (2009), Comparing the Imitative Skills of Children and Nonhuman Apes, in: *Revue de primatology* 1, 1-19

Christensen, W./K. Bicknell/D. McIlwain/J. Sutton (2015), The Sense of Agency and its Role in Strategic Control for Expert Mountain Bikers, in: *Psychology of Consciousness: Theory, Research, and Practice* 2, 340-353

— /J. Sutton/D. J. F. McIlwain (2016), Cognition in Skilled Action: Meshed Control and the Varieties of Skill Experience, in: *Mind and Language* 31, 37-66

Davids, K./P. Glazier/D. Araujo/R. Bartlett (2003), Movement Systems as Dynamical Systems: The Functional Role of Variability and its Implications for Sports Medicine, in: *Sports Medicine* 33, 245-260

Huang, Z./H. Davis IV/A. Wolff/G. Northoff (2017), Thalamo-sensorimotor Functional Connectivity Correlates with World Ranking of Olympic, Elite, and High Performance Athletes, in: *Neural Plasticity*, Article ID 1473783, 1-12

---

than could have made or learned individually, could well be intimately intertwined, even before the emergence of modern humans.

Hull, C. (2020), Prediction Signals in the Cerebellum: Beyond Supervised Motor Learning, in: *eLife*, 9, e54073

Krakauer, J. W./P. Mazzoni (2011), Human Sensorimotor Learning: Adaptation, Skill, and Beyond, in: *Current Opinion in Neurobiology* 21, 636-644

Musgrave, S./E. Lonsdorf/D. Morgan/C. Sanz (2021), The Ontogeny of Termite Gathering Among Chimpanzees in the Goualougo Triangle, Republic of Congo, in: *American Journal of Physical Anthropology* 174, 187-200

Pargetter, J./N. Khreisheh/D. Stout (2019), Understanding Stone Tool-making Skill Acquisition: Experimental Methods and Evolutionary Implications, in: *Journal of Human Evolution* 133, 146-166

Railton, P. (2017), At the Core of Our Capacity to Act for a Reason: The Affective System and Evaluative Model-based Learning and Control, in: *Emotion Review* 9, 335-342

Schmidt, M. F. H./J. A. Sommerville (2011), Fairness Expectations and Altruistic Sharing in 15-Month-old Human Infants, in: *PLoS One* 6, e23223

Smetana, J. G./W. M. Rote/M. Jambon/M. Tasopoulos-Chan/M. Villalobos/J. Comer (2012), Developmental Changes and Individual Differences in Young Children's Moral Judgments, in: *Child Development* 83, 683-696

Tomasello, M. (2016), *A Natural History of Human Morality,* Cambridge/MA

Turiel, E. (2006), Thought, Emotions, and Social Interactional Processes in Moral Development, in: M. Klein/J. G. Smetana (eds.), *Handbook of Moral Development*, Mahwah, NJ, 7-35

Van Overwalle, F./K. Baetens/P. Marien/M. Vanderkerckhove (2014), Social Cognition and the Cerebellum: A Meta-analysis of Over 350 fMRI Studies, in: *NeuroImage* 86, 554-572

Warneken, F./M. Tomasello (2006), Altruistic Helping in Human Infants and Young Chimpanzees, in: *Science* 311, 1301-1303

Wellman, H. W. (2014), *Making Minds: How Theory of Mind Develops*, New York