Beyond Automaticity: The Psychological Complexity of Skill

Elisabeth Pacherie, Myrto Mylopoulos

► To cite this version:

Elisabeth Pacherie, Myrto Mylopoulos. Beyond Automaticity: The Psychological Complexity of Skill. *Topoi*, 2021, 40 (3), pp.649-662. 10.1007/s11245-020-09715-0 . ijn_03496645

HAL Id: ijn_03496645

https://hal.science/ijn_03496645

Submitted on 3 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Preprint version. For purposes of quotation please consult the published version:
Pacherie, E., & Mylopoulos, M. (2021). Beyond automaticity: The psychological complexity of skill. *Topoi*, 40(3), 649-662.

Beyond automaticity: The psychological complexity of skill

Elisabeth Pacherie

Institut Jean Nicod, Département d'études cognitives, ENS, EHESS, CNRS, PSL Research
University, Paris, France

Myrto Mylopoulos

Department of Philosophy & Institute of Cognitive Science, Carleton University

Abstract:

The objective of this paper is to characterize the rich interplay between automatic and cognitive control processes that we propose is the hallmark of skill, in contrast to habit, and what accounts for its flexibility. We argue that this interplay isn't entirely hierarchical and static, but rather heterarchical and dynamic. We further argue that it crucially depends on the acquisition of detailed and well-structured action representations and internal models, as well as the concomitant development of metacontrol processes that can be used to shape and balance it.

Keywords: skill; automatic control; cognitive control; structured action representations; internal models; metacontrol.

1. Introduction

In June 2019, Rafael Nadal won the French Tennis Open, establishing a new all-time record of twelve victories at this tournament. Many consider him the greatest clay court tennis player of all time, and all agree that he is a highly skilled athlete. But what exactly does this skill consist in? To be sure, without speed, strength, and general physical fitness, Nadal wouldn't be the champion he is. Clearly, though, these qualities are not sufficient. One could be superlatively fit and still not know how to play tennis, let alone be skilled at it. What more is needed to make one a skilled tennis player?

Most obviously, acquiring a skill requires (lots of) practice. As Haith and Krakauer (2018) emphasize, practice leads to an improvement in the ability to rapidly select appropriate actions, and to execute them accurately, while at the same time leading to greater efficiency, i.e., a reduction of the cognitive load required to perform the task. These properties are not unique to skills, however. They are also characteristic of the formation of (good) habits.

What is it then that sets apart skilled from habitual action? Philosophers, when they pay heed to the distinction, are apt to be blunt, contrasting the intelligence of skilled action with the unintelligent or mechanical character of habitual action. Psychologists, in turn, commonly characterize habitual actions as automatic, inflexible, and stimulus-driven; by contrast, they insist that although there are elements of automaticity in skill, skilled actions remain goal-directed and highly flexible.

It is central, then, to a full understanding of skilled action to give some account of the intelligence that it exhibits. While different characterizations of what this intelligence amounts to are proposed, we follow others (e.g., Levy 2017) in taking flexibility to be at the heart of this phenomenon. Skilled behavior is highly sensitive to the goals of the agent and the nuances of a given action context. The skilled agent is capable of responding appropriately to novel or unusual

situations. What the skilled action theorist must therefore explain is how this flexibility gets trained up and what processes and representations it involves. A characterization of skilled actions as sharing some of the properties of habits, such as speed, accuracy and efficiency, while at the same time remaining goal-directed and flexible, suggests that skills involve a particular form of interplay between automatic and cognitive control processes.

The objective of this paper is to characterize the interplay between automatic and cognitive control processes that we propose is the hallmark of skill and what accounts for its flexibility. We will argue that this interplay isn't entirely hierarchical and static but rather heterarchical and dynamic. We will further argue that such an interplay depends (i) on the acquisition of detailed and well-structured action representations and internal models, and (ii) on the concomitant development of metacontrol processes that can be used to shape and balance it effectively. Note that while we think that the account of skilled action we propose to articulate could in principle be extended to skilled mental action, our focus here will be on overt bodily action.

In section 2, we start by fleshing out the distinction between habits and skills and we offer a brief general characterization of the intelligence of skill understood as flexibility. In section 3, we draw on the psychological literature to articulate and nuance the distinction between automatic and cognitive control processes; in particular, we discuss gradations in automatic and cognitive control. In section 4, we distinguish between a hierarchical and a heterarchical view of the interplay between these control processes and look at some existing hybrid accounts of skilled action and outline the common commitments of those that take the heterarchical approach. We point out that one limitation of these accounts is that they say little about the types of psychological representations that experts deploy, so they are at the very least incomplete. In section 5, we argue that what distinguishes highly skilled from less skilled performers and novices (e.g., the professional tennis player from the amateur or the beginner) is the structural complexity and accuracy of their mental representations in their domain of expertise. We introduce a distinction between two types of mental representation in experts, representations that support action execution and representations that support action selection. Finally, in section 6 we offer a sketch of how an expert performer comes to have the capacity for what we call *metacontrol*, that is, the ability to successfully arbitrate between more cognitive and more automatic modes of control with respect to both action selection and implementation. We propose that metacognition plays a central role in this process and we close by suggesting that modeling metacontrol in this way may give us a natural way of explaining the distinctive phenomenology of skill.

2. Skills vs. habits

As noted by Douskos (2017), contemporary analytic philosophers often overlook the distinction between habits and skills, taking it to be of little philosophical significance or sometimes not even acknowledging it and using the terms "habit" and "skill" interchangeably. Their indifference to the distinction may be excused in part by the fact that it is often blurred in our ordinary discourse. We talk of Rafael Nadal's skill as a tennis player, but we also talk of tying shoelaces or balancing on a bicycle as skills that children must acquire. Yet, once these latter "skills" have been mastered, their performance is typically a matter of routine. In contrast, saying that Nadal plays tennis routinely would amount to slander. The difference between habits and skills certainly becomes more vivid when instead of contrasting habits with everyday skills (e.g. tying one's shoelaces), we contrast them with the kinds of expert skills demonstrated by Nadal as a tennis player or Martha Argerich as a pianist.

In this section, we begin by outlining four characteristics that help to ground the distinction between habit and skill. Our discussion is guided by the converging accounts of this distinction offered by Ryle (1949) and, more recently, Annas (2011), who are among the minority of philosophers to explicitly address it. In addition to helping demarcate habit from skill, these characteristics also serve as explananda for a complete theory of skill to address.

Firstly, both Ryle (1949) and Annas (2011) contrast the rigidity and uniformity of habitual behavior with the flexibility and variance of skill. Thus, according to Ryle: "It is of the essence of merely habitual practices that one performance is a replica of its predecessors. It is of the essence of intelligent practices that one performance is modified by its predecessors" (1949/2009: 30). Exercises of "merely habitual practices" are "uniform", while exercises of skill are "indefinitely-heterogeneous" (1949/2009: 32; see also Annas 2011: 102). Indeed, this variability in the situations that a skilled agent faces, and in the types of responses that the agent must select and execute in order to cope with this variability, has not gone unnoticed by elite performers. Here is Nadal reflecting on this very point:

"You might think that after millions and millions of balls I've hit, I'd have the basic shots of tennis show up, that reliably hitting a true, smooth clean shot every time would be a piece of cake. But it isn't. Not just because every day you wake up feeling differently, but because *every shot is different; every single one*. From the moment the ball is in motion, it comes at you at an infinitesimal [sic] number of angles and speeds, with more topspin, or backspin, or flatter or higher. The differences might be minute, microscopic, but so are the variations your body makes—shoulders, elbows, wrists, hips, ankles, knees—in every shot. And there are so many other factors—the weather, the surface, the rival. No ball

arrives the same as another; no shot is identical. So every time you line up to hit a shot, you have to make a split-second judgment as to the trajectory and speed of the ball and then make a split-second decision as to how, how hard, and where you must try to hit the shot back.” (Nadal & Carlin 2011, p.6, emphasis ours)¹.

Secondly, in addition to the flexibility and variance involved in skill, the mindlessness of habits should be contrasted with the thoughtful character of exercises of skill. Taking expert piano playing as her example, Annas elaborates on this point, writing that “the expertise is not detached from the person’s ability to think and decide consciously; the playing is continually responsive to my thought about the piece, my decisions to speed up or slow down, and the like” (2011: 102; see also Ryle 1949/2009: 30). By contrast, when it comes to habitual behavior, such as setting an alarm clock before bed or folding the laundry, there is rarely any thought or decision involved.

Thirdly, it is important to contrast the modes of acquisition of habits versus skills. Habits are acquired by drill and repetition and are largely the result of conditioning, while skill acquisition requires training. Training may incorporate a fair share of drill and repetition but it goes beyond this: “It involves the stimulation by criticism and example of the pupil’s own judgment. He learns how to do things thinking what he is doing, so that every operation performed is itself a new lesson to him how to perform better.” (Ryle 1949/2009: 30-31).

This last point deserves emphasis and leads us to a fourth way of marking out skill from habit: Successful learners strive to improve their skills and to master their craft.² Indeed, the idea that greater expertise is linked to heavy investment in the continuous improvement and refinement of skills is corroborated by empirical studies. For instance, in an in-depth psychosocial study comparing elite athletes, who had competed in the Olympics and in world championships, but not won medals, and super-elite athletes who had competed in these events but in addition had won medals, Hardy et al. (2017) found that the main differences between the two groups were as follows: while super-elite athletes tended to have a dual focus on both the outcomes of their performance as well as mastery, elite athletes tended to focus only on outcomes. In addition, super-elite athletes exhibited a greater tendency towards obsessiveness and perfectionism with regard to both kinds of goals. By contrast, it’s unusual for one to strive to improve the execution

---

¹ We thank John Toner for drawing our attention to this quote from Nadal.

² For recent discussions of the connection between skill and practice, see Montero (2016) and Fridland (2019).

of their habits, for example, folding the laundry more quickly. If anything, pursuing such a goal would invite treating the habit as more akin to a skill.

To recap, drawing from the discussions of both Ryle and Annas, we have seen that habit can be distinguished from skill primarily in terms of the latter's flexibility and variance, which is in turn linked to its thoughtful character, and acquisition by way of training with the aim of improvement. So it would seem that in order to give a full account of the specific nature of skill and what sets it apart from habit, we must give some account of how skilled behavior comes to be flexible in this particular way and what underlying representational structures and control processes make such flexibility possible.

We begin this task by offering a working characterization of flexibility. In our view, flexibility comes in degrees, and it is a matter of sensitivity to different aspects of action control across a range of circumstances.³ More specifically, there are at least three dimensions of flexibility that need to be accounted for. Firstly, flexibility involves sensitivity to immediate sensory changes in the environment (*sensory sensitivity*). The skilled performer does not act in a fixed, rigid manner, but rather adaptively responds to minute differences in the environment to ensure optimal performance as an action unfolds. As the quote from Nadal above makes clear, no two shots are identical. The skilled tennis player responds to subtle variations in speed, angle or spin of the ball by way of delicate adjustments in bodily posture and joint angles.

Secondly, flexibility requires sensitivity to the situation of action as a whole (*situational sensitivity*). This involves attending to and responding to higher-level aspects of the situation, namely, perceived properties of the context or environment that are relevant to an agent's goal. To get clearer on the difference between these two types of sensitivity, let us consider the difference between novice and skilled hockey players. In contrast to novices, skilled hockey players are not only sensitive to features like the spatial location of the puck and their grip on the hockey stick, but also situational features relevant to higher-order goals like defensive advantages afforded by certain positions on the ice and recognition of particular plays the opposing team is running (Schmidt & Lee 2014). So skilled experts are able to recognise and process the meaning of perceptual features of a situation in ways that are not exhausted by their sensitivity to basic

---

³ Here we follow Shepherd (2014) in thinking that, in general, "[a]n agent in control is poised to handle any number of extenuating circumstances" (p.399), so that a full understanding of the degree of control that an agent possesses will be evaluated in terms of the agent's sensitivity to various features of these circumstances.

features like the size and position of objects in their milieu, and that afford them a second type of flexibility with respect to their environment that involves more complex forms of processing.

Thirdly, flexibility requires that motor processes are sensitive to the agent's personal-level intentions and goals (*strategic sensitivity*). We do not wish to claim here that this sensitivity must be entirely unmediated; in fact as we will see later, on our own account, this sensitivity is enabled by way of intermediate action representations and attentional processes. But this top-down sensitivity is an important component of skilled action, as it allows the expert performer to control and adjust their movements in ways corresponding to the detail and specificity present in the contents of their goal states.

To illustrate this last dimension of flexibility, consider that at the strategic level, based on his initial assessment of his own strengths and weaknesses and those of his current opponent, Nadal may form the intention to play aggressively from the back of the court. Had he played a different opponent he might have chosen a different strategy. The adoption of a given strategy will influence the way in which he exploits situations during the game and shape how he deploys situational awareness. If, for instance, at some point of the game, the respective current positions of both players and the direction and speed of the ball moving towards him, give him the opportunity to either perform a drop shot or a deep down-the-line backhand shot, his strategic intention will make the selection of the latter shot more likely. Conversely, strategic sensitivity also means that one should be able to revise one's strategic intentions to ensure one's goals are met. Suppose, for instance, that Nadal realizes during the play that his opponent is a tougher baseline defender than he had anticipated, he may decide to change his strategy and force his opponent to step to the net more often.

Having sketched this general picture of flexibility, we now turn to discuss the nature of the different types of control processes that underlie it, in particular cognitive and automatic control processes, and the ways in which they interact.

3. Automatic vs. cognitive control processes

Classical views of automatic vs. cognitive control processes (e.g. Posner and Snyder 1975; Shiffrin and Schneider 1977) often take a monolithic, all or none approach, combining "the idea of perfectly correlated features of each processing mode with the idea that both modes are mutually exclusive and that they exhaust the universe of possible processes" (Moors and de

Houwer 2006: 299). Such views analyze automatic processes in terms of a cluster of diagnostic features and contrast them sharply with cognitive processes, which are thought to possess a different set of features. More specifically, automatic processes are familiarly characterized as being goal independent, unconscious, efficient, and fast, whereas cognitive processes are described as goal-directed, conscious, non-efficient, and slow.⁴ Thus, on such views, automatic and cognitive control processes stand on either side of a strict dichotomy, and whether or not a given process is automatic or control depends on its possession of a specified set of features.

However, this once popular view has been increasingly challenged. Cognitive and automatic control processes are now more commonly viewed as processes that come in degrees, in at least three different respects. First, many of the features associated with both cognitive and automatic control processes are best understood as features that are *scalar* rather than *discrete*. Take, for instance, speed and efficiency. A process is efficient to the degree that it does not demand significant cognitive resources, such as attention or working memory. Speed, in turn, is often considered a function of such efficiency, since the deployment of such resources has a time cost. But then it makes little sense to talk of speed or efficiency in absolute terms, since a process can use up more or less working memory or attention, and subsequently take longer or shorter periods of time to run through completion.

Second, some features associated with automatic and cognitive control processes may be *partial vs. complete* relative to whether they correspond to the inputs, outputs, or mediating computations of those processes. Take consciousness for instance. One might have only limited awareness of the elements and steps of a computational control process, while being fully aware of the inputs to the process or the outputs it yields. This might be the case, for instance with respect to the processing of garden path sentences. Though we are not aware of the computations that pertain to the syntactic and semantic parsing of such sentences, we are aware of the inputs to that process, and the output of the process signaling that an incorrect interpretation has taken place.

Third, as many studies have now shown, the features thought to be diagnostic of automatic and cognitive processes are far from being perfectly correlated, to the point that Bargh (1992) claims that one could come up with an example fitting any random combination of automatic and cognitive features (for discussions of these empirical findings, see Moors and de

---

⁴ For reviews of theories of automaticity and detailed discussions of these features and other features that have been linked to automaticity, see Moors and de Houwer (2006) and Moors (2016).

Houwer 2006 and Moors 2016). Less radically, many now advocate a scalar view that takes automatic and cognitive processes to form a continuum rather than a dichotomy (see especially the work of Fridland 2015). They propose that although no unique combination of features may be necessary and sufficient for a process to be either automatic or cognitive, processes could still be assessed as more or less automatic or cognitive depending on how many of the relevant features they possess and to what degree (e.g., Ledoux & Daw 2018).

Finally, many theorists now recognize that most actions cannot be identified as involving cognitive processes only or automatic processes only. Rather they involve a mix of cognitive and automatic processes and interactions between them. There has been some resistance to this last claim in the context of skill. For example, both Fitts and Posner's 1967 model of skill acquisition in psychology and Dreyfus & Dreyfus's (1986) model in philosophy view skill as a progression from rule-following and significant cognitive control at the early novice stages to an expert stage wherein skilled performance becomes heavily automated, movements exhibit maximal accuracy, consistency, and efficiency, and little or no cognitive control is required except in novel or unusual circumstances. Thus, Dreyfus and Dreyfus urge, "when things are proceeding normally, experts don't solve problems and don't make decisions, they do what normally works" (p.31).

More recently, however, there is growing recognition that, though the move from novice to expert does involve increased automaticity, it does so in a way that involves the "freeing up" of cognitive resources that then allows for a different mode of engagement via cognitive control processes. As Christensen et al. (2016) put it, a theory of skill must then "address not only the automation of aspects of control but also the shift in the role of cognitive control, its main processes, and the relations between automatic and cognitive control in performance" (p.38).

How should we conceive of this interplay between cognitive and automatic control processes and their underlying representations? We address these questions in the following sections.

4. Hierarchical vs. Heterarchical Control

It has been proposed (Pacherie 2008; Mylopoulos & Pacherie 2018) that the cognitive architecture underpinning intentional action involves a hierarchy of representations and processes that specify and control the action in a gradually more fine-grained way. This model (the DPM model) distinguishes three main stages in the process of action specification, each corresponding to a different level or layer of intention, with intentions at each level having a

distinctive role to play in the guidance and monitoring of the action. It thus proposes a threefold distinction among distal, proximal, and motor intentions (D-intentions, P-intentions and M-intentions for short), where D-intentions are more abstract, conceptual representations of action goals integrated with the agent's beliefs and desires, P-intentions are more concrete, contextualized representations of goals that integrate perceptual information about the situation of action, and M-intentions encode action goals in a motoric format directly suitable for action execution. In addition, the model takes it that intentions at each of the three levels have a distinctive role to play in the control and monitoring of ongoing actions, with D-intentions being responsible for strategic control, P-intentions for situational control, and M-intentions for motor control.

One may be tempted to think that the goal selection and control processes at the top of this hierarchy are cognitive in nature and that the further down the hierarchy one goes, the more automatic they become, with strategic control involving only cognitive control processes and motor control only automatic processes (e.g., Dezfouli & Baleine 2013). On such a view, cognitive and automatic control processes would themselves be hierarchically organized in a way that corresponds directly to the action representation hierarchy. Alternatively, one might think that the cognitive-automatic control gradient need not be aligned with the action representation hierarchy and that the interplay of cognitive and automatic control processes can be more flexible and *heterarchical*. In other words, a representational hierarchy is in principle compatible with a control heterarchy. For instance, Cushman and Morris (2015), from whom we borrow the expression “heterarchical control”, propose that whether or not one relies on automated control processes (what they call “habitual control”) or on more cognitive control processes (what they call “planned control”) is not necessarily a matter of the level of abstraction in the hierarchy of action specification. Rather, this depends on the amount of variability present at a given level in terms of the action types that are available to serve the corresponding goals within the given context and on the amount of experience the agent has with the task as specified at that level. They illustrate their point by means of the following example:

Consider, for instance, a seasoned journalist who reports on new events each day. At a high level of abstraction, her reporting is structured around a repetitive series of goal-directed actions: follow leads, interview sources, evade meddling editors, etc. Because these actions are reliably valuable for any news event, their selection is an excellent candidate for habitual control. The concrete steps necessary to carry out any individual action will be highly variable, however—optimal behavior when interviewing a pop star

may be suboptimal when interviewing the Pope. Thus, the implementation of the abstract actions is an excellent candidate for planning. (2015: 13817)

In other words, Cushman and Morris (2015) propose that humans will flexibly adapt the use of cognitive and automatic control across levels of hierarchically organized behavior to suit the demands of the task at hand. In effect such a proposal involves combining the idea that humans mentally organize their behavior around hierarchically structured goals and subgoals (representational hierarchy) with the idea of a heterarchy of behavioral control.

Indeed, this heterarchical model of action control is reflected to varying degrees in some recent 'hybrid' accounts of skill that aim to specify the respective roles of cognitive and automatic processes in skilled, flexible performance (for examples of hybrid accounts, see Christensen et al. 2015, 2016; Fridland 2014, 2015, 2017; Levy 2017; Montero 2016; Papineau 2013, 2015; Shepherd 2019). Some of these accounts allow for significant contributions to flexibility from a range of cognitive control processes and even a direct influence of such processes on motor implementation and execution, while others take a more narrow view of cognitive control and allow for only limited contributions at the motor level.

As an instance of the latter, consider Papineau's (2015) hybrid view of skill. He allows that cognitive control, which he restricts to deliberate thought and focus, and refers to as "keeping one's mind right", is of utmost importance for expert performance, even when significant automatization of motor routines has been achieved. He holds that this form of cognitive control must not be directed at detailed motor routines themselves (e.g., adjusting one's hold on the tennis racket in a certain way), but rather at the level of what the agent intends to do (e.g., return the opponent's serve). On this view, motor implementation is only indirectly shaped by cognitive control in that it triggers one set of motor routines over another, depending on what the agent intends to do. The rest occurs automatically by way of heavily conditioned reflex-like responses.

By contrast, consider the hybrid account that has recently been defended by Christensen et al. (2016), which presents a more variegated view of cognitive control, and is more strongly heterarchical. On this account, cognitive control is associated with the operation of an executive system with a range of primary functions including "controlling attention, the active maintenance and processing of information (working memory), the flexible integration of information related to the current situation and activities, setting and switching between goals, establishing an action or task 'set' (a processing configuration for the situation), inhibiting inappropriate responses, forming action plans, decision-making, and problem solving" (2016: 40). Thus, it is not restricted to the type of conscious deliberate reasoning processes that are often associated with this term.

Christensen et al. also defend a nuanced understanding of automaticity, allowing for a more graded notion, with weaker forms of automaticity possessing some but not all of the features traditionally associated with it.

Importantly, and in line with a heterarchical view of action control, Christensen et al. (2016) recognize rich systemic interactions between cognitive and automatic processes. They maintain that as the motor routines in skilled action become progressively more automatized, cognitive control is diminished, but still makes important contributions and in many cases directly influences the implementation of motor routines (p.43). These influences are not restricted to situations where something has gone wrong with the action and it needs appropriate adjustments, nor to novel circumstances that do not as easily allow the agent to fall back on automatized routines. Rather, Christensen et al. argue that at advanced stages of skill, even when a situation is familiar to the agent (e.g., returning an opponent's volley), this does not mean that it is easy. Rather, it will likely involve a level of difficulty that is sufficiently demanding for the agent such that cognitive control has a significant role to play.

They further emphasize that cognitive control is not just concerned with high-level, relatively coarse-grained strategic aspects of action performance, such as, for instance, decisions to play defensively or offensively, but also with what, following Endsley (1995, 2006) they call *situation awareness*. As characterized by Endsley, situation awareness involves the moment by moment perception of relevant elements from the environment, the comprehension of their meaning, and the ability to project from current events and dynamics to anticipated events in the near future. Christensen et al. (2016) point out that situation awareness depends on the flexible integration of information concerning the current situation, which is one of the main functions of cognitive control. Situation awareness can involve explicit inferential reasoning processes, but also often occurs without such processes and is closely linked to attentional control. It is typically constructed progressively and revised over time. Compared to representations of strategic aspects of action performance, situation awareness involves much more fine-grained and time-sensitive representations of the situation at hand and can thus serve to establish a cognitive and motor configuration appropriate to the context that directly influences action execution.

Similarly, Fridland's (2014, 2017) hybrid account of skill is heterarchical in that it stresses the role of strategic goal states in the initial programming of the motor system, and also insists that the interaction continues beyond this stage, and on through the action's completion. She writes, "the best evidence we have indicates that fine-grained, automatic motor processes instantiated in motor skills are not simply causally connected to intentional states but, rather,

continue to be semantically sensitive and responsive to personal-level goals throughout execution" (2017, p.1557). She suggests that this ongoing sensitivity is served via selective top-down attentional processes that are directed to features of the environment relevant to the implementation of the action goals specified by intentions at the level of strategic control. Finally, she proposes that although motor control tends to be heavily automatized by way of practice and training and is not guided by strategic intentions each step of the way, intentions can nevertheless intervene at various junctions if necessary.

The hybrid accounts of skilled action proposed by Fridland and Christensen et al. thus operate with a more subtle distinction between automatic and cognitive processes and allow for a less hierarchical, more flexible interplay between them. Both accounts posit the need for an intermediate form of control between high-level cognitive control by intentions and largely automatic motor control to account for the flexibility of skilled action. Both accounts offer suggestions regarding the processes at work at this intermediate level, selective attention according to Fridland and information integration processes according to Christensen et al. Both also take these processes to possess some of the marks traditionally associated with automaticity and some of the marks associated with cognitive processes. In all these ways our accounts are similar, but our proposal differs from those of Fridland and Christensen et al. in two main ways.

First, our account is more strongly heterarchical in that both Christensen et al. (2016) and Fridland (2014) maintain that at higher levels of control there is little or no automation. Thus, Fridland tells us that an "... important feature of strategic control is that it does not automate. That is, while many aspects of skill are learned over time and become automatized into routines that can be run in the absence of explicit attention, strategic control does not" (p.2745). And Christensen et al. (2016) write that, "situation control and higher strategic control do not tend to automate strongly" (p.49). On our view, however, even these higher levels of control can be automatic.

Second, neither Fridland (2014) nor Christensen et al.'s (2016) account tells us much about how the attentional and information-integration processes they posit are themselves guided and how precisely they contribute to the formation of appropriate motor routines. For instance, Fridland writes that "the trained expert need not voluntarily deploy any particular motor routine when executing a skill. Rather, the right motor routines are deployed as a result of the agent initiating her skilled action together with her capacity to selectively attend" (2748). Similarly, Christensen et al. (2016) tell us that "as situation interpretation develops, attention is directed to relevant information which serves to elaborate or revise the interpretation. Situation awareness

serves to establish a cognitive and motor configuration appropriate to the context” (p.43). But the question of how precisely selective attention yields the “right” or “appropriate” motor routines and configurations remains open.

Our contention is that in order to answer this question one must have some account of the precise *representations* underlying this *intermediate level* of control. Fridland does not offer us such an account. Christensen et al. introduce us to the notion of an action ‘gist’, which is a goal representation more detailed than strategic intention but not as detailed as motor representations, “specifying not just an action type but also a particular way of performing the action appropriate to the circumstance” (p.43) and thus directly “shaping” and “regulating” action execution. But the precise nature and structure of these action gists remain underspecified.

So while we are highly sympathetic to these approaches and their strategy for accounting for the flexibility of skilled action, we also think that more needs to be said about the specific knowledge structures and mental representations that subtend expert action.⁵

5. The Action Representations of Experts

In contrast to classical theories of skill acquisition, that saw it as a process of automatization (e.g., Fitts and Posner 1967), Ericsson (2006) argues that rather than being the mark of expert performance, full automatization corresponds to *arrested development*. In his view, “expert performers counteract automaticity by developing increasingly complex mental representations to attain higher levels of control over their performance” (2006: 685). Thus, what distinguishes expert performers from novices or merely competent performers (e.g., the professional tennis player from the amateur or beginner) is, at least in part, the complexity and accuracy of their mental representations of their domain of expertise and action repertoires as well as the way these representations guide action selection and control action implementation.

Here, we distinguish between two types of mental representation that are characteristic of expert action: what we call structured action representations (SARs for short), that is, representations of actions that form part of the motor repertoire in the relevant skill domain (e.g.,

---

⁵ It is also worth noting here that, insofar as our account of the intelligence of skill depends heavily on the types of action representations deployed in expert skill, we depart from views of this intelligence defended by theorists like Dreyfus (2002), who hold that it can be “described and explained without recourse to mind and brain representations” (p. 367)

a slice forehand in tennis or an upward glissando on the piano) and support action implementation and, at higher-levels, internal ⁶models of the domain of skill that support action selection.⁷

5.1 SARs for Action Implementation

As an illustration of SARs, one can consider Bläsing et al.'s (2009) study of the mental representations of actions in the classical ballet repertoire, specifically *Pirouette en dehors* and *Pas assemblé*, in dancers with different skill-levels. Two main considerations guided the choice of these two movements. First, they are highly familiar, both visually and motorically, among professional and amateur dancers, and ubiquitous even in beginners' classes. Second, they present a sufficiently high degree of complexity: both involve several functional phases and require important coordination, but there exist crucial differences in biomechanical structure, with *Pirouette* considered more difficult and involving a rotational movement. A preliminary phase in the study involved breaking each movement into basic units, defined and labelled with the help of movement descriptions from experts and standard references on classical dance training. For instance, *Pas assemblé* was decomposed into 9 basic units: (1) stand, left foot in front; (2) bend knees; (3) right foot slides to side, (4) lift right leg, (5) jump from left leg, (6) stretch left leg in air; (7) join legs; (8) land on both feet, (9) bend knees, stretch, with (1)-(2) correspond to the preparatory tension building functional phase, (3)-(7) to the jumping functional phase and (8)-(9) to the landing phase.

In their study, Bläsing et al. used an analytic method called Structured Dimensional Analysis-Motorics (Schack, 2001, 2002), based on a method well established in the field of cognitive psychology to study relational structure in a given set of concepts and adapted for the analysis of movements. As a first step, participants were asked to perform a multiple sorting task where each unit in turn served as an anchor and where they had to decide for every other unit whether it was “functionally close” to the anchor while performing the movements, thus providing information on the representational distance between units. Based on this information, a hierarchical cluster analysis was then used to transform this set of units into a hierarchical

---

⁶ Importantly, these are not the same as the forward and inverse models often posited in the motor control literature, though these are referred to as internal models as well. For a discussion of these models, see Wolpert & Kawato (1998).

⁷ Note that the two types of representations can come apart. For instance, a sport commentator would be expected to have good internal models of the domains of, say, tennis or soccer, but not necessarily adequate SARs for actions in these domains. Her job is to comment intelligently on games, not to expertly play tennis or soccer. A professional player in contrast is supposed to have acquired both types of representations.

structure. Third, a factor analysis was used to reveal the dimensions in this structured set of units, and finally, the cluster solutions were tested for invariance within- and between-groups.

Bläsing et al. found that the experts' representational structures for *Pirouette en dehors* and *Pas assemblé* were different from those of amateurs and novices. The structure of the representations in experts had a distinctive hierarchical organization that was well matched with the functional and biomechanical demands of the action, and exhibited very little variation from dancer to dancer. In contrast, in amateurs, action representations showed less hierarchical organization, were less well matched with the functional phases of the action, and their inter-individual variability was much higher. In novices, strikingly, hierarchical structure was almost absent.

SARs are both representational and control structures. A well-formed SAR is a hierarchically organized assemblage of motor schemas. Drawing from influential work in the motor control literature, motor schemas can be understood as involving representations that encode the general form or pattern of an action, that is, the features of a motor act that remain constant across its various instances (Schmidt 1975, 2003; Arbib 2003; Jeannerod 1997). These invariant features correspond to the “deep structure” of an action and include aspects such as the functional organization of events, their order and relative timing, as well as their general spatial configuration. For instance, as we have seen earlier, *Pas assemblé* is composed of three main functional phases, tension build-up, jump and landing, with each phase itself involving sub-phases and ultimately individual motor schemas.

In contrast, more superficial aspects of the movements will change from one occasion to the next. These variable or surface features of a motor act correspond to the parameters of a motor schema. On different occasions, the parameters will take different values determining how the motor schema will be expressed. The surface features that need to be specified each time the movement is performed include elements such as the movement speed, its amplitude, its direction, or the effector used. So, for instance, how high one jumps or how quickly one stretches one's left leg in the air and in what direction in a given performance of the *Pas assemblé* are its variable features.

In order to determine how an action should be performed on a given occasion, parameter values adapted to the situation at hand must be specified. Schema theory proposes that information about initial conditions, parameter values, action outcomes, and sensory feedback are stored together every time a movement is made and that, as an agent learns through practice, mapping rules are abstracted from this stored information across a number of performance

instances. For instance, after sufficient practice throwing balls at different distances, a mapping rule relating parameter values to outcome distances will have been acquired. The agent will then be able to rely on the schema to set the parameter value for a throw at a specific distance, even if he or she has never before produced a throw at that exact distance. Thus, in addition to the general form of an action, what the agent learns through practice are sets of mapping rules relating initial conditions, parameter values, action outcomes and sensory feedback, such that motor schemas can become more and more fine-tuned over time and be successfully deployed even when novel conditions are encountered.

Importantly, another result of practice is the “chunking” of motor schemas into SARs that serve to integrate them into a single functionally and temporally organized representational unit corresponding to a complex action type. Thus, for instance, through a process of chunking, motor schemas corresponding to a sequence of movements [A][B][C][D][E][F] may become integrated into two SARs corresponding to [[A] [B] [C]] and [[D] [E] [F]] (see Fridland 2019 and Pavese 2017 for important discussions of this process). Some motor schemas are elementary or basic in the sense that they correspond to action types that cannot be “broken down” any further into subgoals. When elementary motor schemas are chunked together, a SAR is born. But this process of chunking can iterate, such that SARs can themselves be further integrated into yet more complex SARs. So, for instance, the SARs corresponding to action types [A, B, C] and [D, E, F] could themselves be chunked together into a single SAR with the structure [[[A] [B] [C]] [[D] [E] [F]]].

The SARs resulting from chunking retain a nested hierarchical structure, such that each chunked motor schema corresponds to a node in the hierarchy. This means that the various nodes in a hierarchically structured action representation can function as intermediary control nodes. In particular, they can control attentional deployment to the relevant environmental properties for the purposes of parameterizing the relevant motor schemas. We do not claim here that in order to play this role these representations must of necessity be available to consciousness or result from deliberation. Indeed, their deployment will typically be automatically triggered by way of the strategic intention that is guiding the relevant action. SARs also enable intentional intervention at various junctures of the action, thus making it possible to exert some form of cognitive control over action implementation. For instance, feint actions in tennis (say, ostensibly preparing a forehand attack shot and at the last moment turning it into a drop shot) depend on the possibility of exerting cognitive control over the action, inhibiting the last functional steps of the attack shot and substituting it for a drop shot. Such an intervention wouldn’t be possible if the motor units composing an action (e.g., attack hit) were simply chained with or concatenated, as once triggered

the whole motor sequence would automatically unfold. SARs can thus support both automatic and cognitive control processes. Intermediary control nodes may either use default values, corresponding to the automatic unfolding of the action once triggered (e.g., a regular attack hit), or use specific non default values when cognitive control is exerted (e.g., a feint attack hit or a attack hit performed in difficult circumstances and requiring extreme precision). The existence of these intermediary control nodes thus allows for greater plasticity than intentions and motor commands can on their own afford the agent.

In addition, as Fridland (2019) emphasizes, skill learning involves deliberate attempts not only to improve one's success at achieving one's ends, but also the *technique or means* by which those ends are achieved. A skilled tennis player will not only aim to improve her top-spin forehand shot, but also her forearm pronation, full shoulder turn, wrist rotation, eye contact, etc. This type of ongoing deliberate practice, considered by Ericsson (2006) as crucial to the development of superior expert performance, helps to establish the specific mappings and hierarchical organisation of the structured action representation, and to further fine-tune them as the agent increases their skill level.

Thus, one mark of expertise lies in the superior quality of the action representations that mediate skilled performance. Well-structured action representations allow performers to select better movements, to better anticipate action effects and future events, and thus to better (and more finely) monitor and control various aspects of their performance. Expertise is not just a matter of how accurately one executes selected actions, however, it is also a matter of selecting appropriate actions in the first place. We now turn to this issue.

5.2 Internal models for action selection

Consider Pierre, a born and bred Parisian. Pierre has lived in various neighborhoods in Paris. He doesn't own a car and he takes public transportation to go from one place to another. He prides himself on knowing the public transportation network inside out. How does he decide which means of transportation to use to get from point A to point B? It depends. If he wants to get from home to work, he can certainly rely on routine, knowing from experience that taking bus 27 is the fastest way to get there. If, in contrast, the route is unfamiliar, say he needs to go from his dentist appointment to the opera ticket office, he will use his internal model of the transportation network, figure out what the options are and compute the best route. One option would be to take bus 52 but he knows that at that time of the day the traffic will be heavy in the districts the bus passes through and it will take an inordinate amount of time. Another option would be to take metro line

9 and then walk for 10 minutes but it's raining and he doesn't have his umbrella with him. Finally, he could take line 6 and then line 8. This involves a transfer at the station La Motte Picquet, but it is rather a short one and this portion of line 6 is an elevated railway with scenic views. Pierre opts for this last option.

Internal models are mental representations of the causal structure of the environment, supporting predictions concerning (potentially probabilistic) action outcomes. Thus, Pierre's internal model of the Parisian public transportation network represents the bus and metro lines, their stops, the points where they connect and various bits of ancillary information (e.g., scenic views, long and smelly metro corridors, slow buses, etc.). Internal models can be used to compute on the fly the optimal action given the structure of the environment and the agent's goal, eg., whether taking route A will get Pierre where he wants to go and whether it will do it more efficiently than route B. Internal models can be more or less detailed. Someone who has lived in Paris for only a couple of years won't have as complete and detailed an internal model of public transportation as an old-timer like Pierre.

Pierre can thus either rely on habit or consult his internal model of the public transportation network to decide which route to take. This echoes the distinction drawn in modern decision research (e.g., Dickinson 1985; Kahneman 2003; Kool et al. 2018; Sloman 1996) between two main strategies for solving an action selection problem: the habitual or model-free strategy and the planning or model-based strategy. The model-free strategy can be seen as exemplifying Thorndike's law of effect (Thorndike 1911), according to which responses to a stimulus that produce a satisfying effect become more strongly linked to the stimulus, and responses that produce a discomforting effect become more weakly linked to it. Roughly, then, a value is associated with an action based on how it has been rewarded in past experience. This value is incrementally updated, and the action with the highest associated value is selected (e.g., taking bus 27 to get from home to work). This strategy is called model-free because it does not rely on an internal model of the environment. Rather, values associated with actions are stored in a cached format and can be rapidly retrieved. On the upside, this decision strategy is computationally frugal, hence fast and efficient. On the downside, it lacks flexibility in the face of change. If there is a change in the environment or the task in a way that warrants different actions being selected, the entire set of cached values will be outdated and will need to be slowly relearned through experience. In contrast, model-based strategies appeal to internal models of the environment. Rather than relying on cached values, they use internal models to compute the optimal action given the agent's goal. Model-based decision making will therefore be more

computationally costly than model-free decision making, but at the same time it will allow for greater flexibility. For instance, internal models can be modified locally to accommodate changes rather than having to be entirely relearned. The two strategies thus represent opposite trade-offs between efficiency and flexibility.

Pierre's example might be somewhat misleading in suggesting that the use of internal models necessarily involves slow, conscious deliberation. If it were so, model-based planning wouldn't be compatible with the temporal constraints of real-time execution typical of many skills, and cognitive control involving such slow cognitive deliberation would only have its place during learning and training phases.⁸ However, while using internal models to compute, say, novel routes, might involve elaborate forms of deliberation and planning, deliberative and planning processes more generally need not be conscious and indeed will often remain implicit and non-conscious. As we emphasized in section 3, cognitive processing shouldn't be equated with slow, conscious processing. As Christensen et al. also stress, "cognitive control, understood in terms of the operation of the executive system, involves a broader range of processes than just conscious reasoning" (2016: 40).

We propose that a second mark distinguishing experts from ordinary mortals is that experts have acquired very detailed and accurate models of their domain at both the strategic and the situational or tactical levels.⁹ Having these models allows them to make optimal decisions, even in unusual and challenging situations. For instance, Nadal knows exactly what his style of play is, what his strengths and weaknesses are and how to make the best of both of them. His internal model of tennis also includes detailed information on other styles of play and their associated strengths and weaknesses and information on how best to exploit the weaknesses of his opponents to his advantage. When playing a particular opponent, he will use his internal model together with information, gathered from previous encounters and from observation of his opponent during warm-up, about the weaknesses in technique, movement, reaction time, as well as strengths, preferred and most consistent types of shots, speed of first and second serve, etc., to decide on how to play against this opponent on this specific court in that particular tournament.

---

⁸ We thank two anonymous reviewers for this journal for pressing this issue.

⁹ Christensen et al. (2015) also emphasize the role played by internal models in the control of complex skilled action. They call them "causal control models", reserving the term "internal model" for the forward and inverse models often discussed in the motor control literature (see fn. 6). In particular, they propose that causal control models "incorporate explicit representations of causal relations", "are at least partly accessible to awareness and participate in high order control" (p. 346). But while they take it that skilled action control involves both more automatic control processes and control processes based on causal control models, they do not explicitly address the metacontrol issue of how agents arbitrate between these two forms of control (see Section 6).

For instance, he may decide on a strategy of playing aggressively from the baseline, hitting the balls on the rise and putting pressure on the opponent. During the game Nadal will also have to make more tactical decisions on how to play a specific point given the current situation. For instance, having noticed that his opponent has lost stamina in the past two or three points, Nadal may decide to take him for a long rally to tire him even further and if he hasn't yet broken down after six shots finish him off with a drop shot close to the net. Thus while SARs guide the implementation of selected shots (say, a dropshot), internal models are used to decide on the strategy and tactics of the game, from the overall strategic decision to play aggressively from the baseline all the way to the tactical decision to win the point with a final dropshot.

However, to say that having detailed internal models of the domain allows an expert to make optimal strategic and tactical decisions isn't to say that they should always rely on their internal models. Similarly, to say that SARs provide the expert with a range of intermediate control nodes over her actions isn't to say that the expert should in each of her actions exert cognitive control over each of these nodes. If they did, the price to pay for flexibility would be loss of efficiency. If Nadal has already played a particular opponent a number of times and always won his matches against him, he doesn't need to think long and hard about strategy, he can simply rely on the same strategy he used in past encounters. If in contrast, his opponent is a new rising star he is playing for the first time, he might want to carefully consider which strategy he should use against him. Thus, we take it, an expert is one who knows when to let things happen (release control to more habitual or automatic processes) and when to make things happen (exercise more cognitive control over action selection and execution). Thus, a third hallmark of control in skill is the ability to dynamically alter the balance between more habitual or automatic and more cognitive modes of action control, be it at the motor, situational or strategic level. We may call this metaflexibility and the form of control that comes with it metacontrol. In the next section, we offer a sketch of how such control might operate with respect to both action selection and implementation.

6. Metacontrol

The first thing to observe is that the psychological capacity that underlies metacontrol is naturally construed as a form of metacognition. In general, metacognition refers to the capacity to monitor and evaluate one's cognitive states and processes for the purposes of controlling and regulating them. Research on metacognition has identified a role for this capacity in a variety of cognitive domains, perhaps the most widely examined being perceptual decision-making and memory

retrieval. Here we wish to suggest that metacognition can itself be directed towards control processes for the purposes of *their* control. Indeed, metacognitive processes have already been implicated in the metacontrol of both action selection and implementation. We take each of these in turn.

Getting clear on the details of how an agent arbitrates between model-based and habitual control strategies in selecting a course of action is currently a central aim of decision-making theory. Several models of this arbitration process have been put forward (see, e.g., Boureau, Sokol-Hessner, & Daw 2015; Gershman, Horvitz, & Tenenbaum 2015; Griffiths, Lieder, & Goodman 2015; Keramati et al. 2011; Pezzulo, Rigoli, & Chersi 2013), but as of yet no consensus has been reached. Still, many of these models share the core assumption that the balance between model-based and habitual control of action selection is determined on the basis of a metacognitive evaluation that takes the form of a cost-benefit analysis that is sensitive to the intrinsically higher costs of deploying model-based control. Given the extra resources and time such control requires, it would make sense for it to be utilized only when the rewards it would be expected to yield for a given task (e.g., greater accuracy) outweigh the built-in costs compared with those predicted for habitual control (see Kool et al. 2018; Shenhav et al. 2013; Gershman et al. 2015; Griffiths et al. 2015).

This assumption leads to the prediction that an agent's use of model-based control will correlate with increased incentives for a given task, since this is when the costs of such control would be most beneficial. In line with this prediction, Kool et al. (2017) found that on a task where the size of the stakes for each trial was randomly selected, and participants were cued with the reward value on each trial, they showed an increased reliance on model-based control and effortful planning on high vs. low-stakes trials. Importantly, this behavior does not seem to be the result of a simple heuristic that dictates that model-based control should be used when stakes are high, since on a separate task where there was no trade-off between cost and reward the same preference was not displayed. As the authors point out, this further suggests "flexible and adaptive integration of costs and benefits" (Kool et al. 2018, p. 164) in the service of metacontrol.

In the context of skill, this proposal is especially attractive. A skilled agent is one who engages in continuous learning. Experts are committed to mastering their craft, and developing it over time through hours and hours of practice and training. Through this process they build up an extensive history of costs and rewards associated with the use of different control strategies in distinct types of circumstances within their domain. We should thus predict that in light of their

dedicated training, experts are able to fine-tune the kind of cost-benefit analysis that is thought to drive the metacontrol of action selection strategies.

Turning now to metacontrol for action implementation, we observe that such control requires knowledge or awareness of what is working and what isn't as an action is unfolding. If an agent has no access to information about the success of their control processes, then they will have no input on the basis of which to guide their *metacontrol* processes and balance the interplay between cognitive and automatic control.

Metacognition has also been implicated as a source of such awareness. A number of empirical findings seem to suggest that metacognitive monitoring signals that pertain to how successfully an action is being implemented drive an agent's judgements of how much control they have over it. In particular, what they seem to be tracking is how *smoothly* the action is being executed, down to its low-level details (see, e.g., Metcalfe & Greene 2007; Wenke et al. 2010; Chambon & Haggard 2012; Chambon, Sidarus, and Haggard 2014). For instance, certain findings seem to suggest that introducing perturbations during the implementation of a movement results in lower judgments of control even if the intended goal of the action is satisfied (e.g., hitting the target) (Metcalfe & Greene 2007), and that subliminally priming a conflicting response just prior to the movement will lead to lower judgments of overall control (Wenke et al. 2010).

On a popular way of modeling this metacognition for action implementation, what is being tracked at the metacognitive level is the degree to which the predicted sensory consequences of an action correspond to the actual sensory feedback (see Wolpert et al. 1995; Chambon et al. 2014). Again, we see that the skilled agent will be particularly adept at this form of metacontrol. In the course of building up their SARs through practice, training, and performance, the skilled agent will have acquired precise mappings between relevant environmental properties, action parameters, and expected sensory consequences of movements. We should expect, then, that the skilled agent is able to more readily detect when their execution of an action is going less smoothly than anticipated, and when they should intervene with the help of cognitive control resources.

We close by offering the suggestion that one key advantage of understanding the capacities underlying metacontrol with respect to both action selection and implementation as forms of metacognition is that it affords us a natural way of accounting for an important marker of skilled action that some theorists, including both Ryle and Annas, overlook as an important datum to be explained: its distinctive phenomenology (for exceptions, see especially Dreyfus & Dreyfus 1986 and Christensen et al. 2016). The subjective experience of engaging in skilled performance

is dramatically different from that of engaging in a habitual routine. In particular, skill is often marked by a type of experience that has been termed “flow”, a feeling of “effortless absorption” in one’s present activity.¹⁰[6]

Research on the experience of flow has been conducted since the 1960s in the context of domains as diverse as chess, rock climbing, and dancing (Csikszentmihalyi 1975/2000). Flow has been characterized in many ways, but following the influential work of Csikszentmihalyi in this area, it can be understood generally as the “subjective experience of engaging just-manageable challenges by tackling a series of goals, continuously processing feedback about progress, and adjusting action based on this feedback” (Nakamura & Csikszentmihalyi 2009, p.90).

We suggest that aspects of this experience may be the direct result of the metacognitive processes that underlie the skilled agent’s arbitration of control processes for both action selection and implementation. Indeed, with respect to action selection, Kurzban et al. (2013) have argued that what they call the subjective feeling of effort is the product of the very metacognitive monitoring mechanisms that calculate the opportunity costs of deploying model-based strategies on the present task (vs. other tasks that could be performed). Likewise, the metacognitive monitoring of action implementation processes has been thought by many to result in feelings of fluency that then feed into the agent’s control judgments for a given performance (see, e.g., Chambon et al. 2014). So it may be, we suggest, that the “effortless absorption” that is characteristic of flow experiences is at least in part the result of the ongoing arbitration that the skilled agent engages in for the purposes of metacontrol, and that this gives us a natural way of accounting for aspects of the phenomenology of skilled action, and its absence in mere habitual behavior.

7. Conclusion

As we saw in section 2, for Ryle and Annas what distinguishes skill from habit is its plasticity and flexibility, where this flexibility is linked to the thoughtful character of exercises of skill as opposed to the mindlessness of habits. In this paper we have offered an account of the types of

---

¹⁰ Note that this emphasis on a proprietary kind of metacognitive phenomenology for skilled action vs. habit is consistent with the phenomenon of “expertise-induced amnesia”, which concerns the inaccessibility to verbal report of detailed aspects of skilled performance, as evidenced by some expert testimony (for reasons to doubt how widespread these reports are, see Bermúdez 2017). The outputs of metacognition pertain directly to the success of the first-order control process, but do not themselves carry any descriptive information about that process beyond this.

representational structures that underlie this flexibility and the specific form of interplay between automatic and cognitive control processes that enable it.

We proposed that crucial to skill is the superior quality of the action representations acquired by experts and responsible for action implementation, what we called SARs, as well as the possession of detailed and accurate models of the domain of skill at both the strategic and situational levels. We take it, following the insights of Ryle and Annas, that variation rather than mere drill and repetition as well as critical attitude and understanding are crucial to the acquisition of skill because they are crucial to the acquisition and fine-tuning of accurate SARs and accurate models of the domain of skill.

We also proposed that having accurate SARs and internal models makes it possible for skilled performers to exert cognitive control over their actions at both the decision and implementation levels and is key to the flexibility and intelligence of skill. However, skilled action is not just flexible, it is also highly efficient. A further mark of expertise is the ability to shape the interplay of automatic and cognitive processes, and flexibly release control to automatic processes or exert cognitive control to guide action selection and implementation. We called this metaflexibility and the type of control associated with it metacontrol, allowing us to alter the balance between automatic and cognitive control. We take it again that the continuous striving to improve highlighted by Ryle and Annas as characteristic of skill is key to reliable metacontrol. Knowing the extent to which to let things happen and the extent to which to make things happen requires experts to have amassed an extraordinary wealth of experience. Finally, we pointed out that there is a specific metacognitive phenomenology linked to metacontrol, suggesting that a final distinction between habit and skill, overlooked by both Ryle and Annas, concerns their dramatically different phenomenology.

Acknowledgements: Elisabeth Pacherie's work was supported by grants ANR-10-LABX-0087 IEC and ANR-10-IDEX-0001-02 PSL from the French Agence Nationale de la Recherche.

References

Annas, J. (2011). Practical expertise. In J. Bengson and M.A. Moffett (eds.), *Knowing how: Essays on knowledge, mind, and action* (pp. 101-112). Oxford: Oxford University Press.

- Arbib, M. A. (2003). Schema theory. *The Handbook of Brain Theory and Neural Networks* (second ed.), MIT Press, Cambridge, MA, pp. 993–998.
- Bargh, J. A. (1992). The ecology of automaticity: Toward establishing the conditions needed to produce automatic processing effects. *American Journal of Psychology*, 105(2), 181-199.
- Bermúdez, J. P. (2017). Do we reflect while performing skillful actions? Automaticity, control, and the perils of distraction. *Philosophical Psychology*, 30(7), 896-924.
- Bläsing, B., Tenenbaum, G., & Schack, T. (2009). The cognitive structure of movements in classical dance. *Psychology of Sport and Exercise*, 10(3), 350-360.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioural and brain sciences*, 18(2), 227-247.
- Burnston, D. (2017). Interface problems in the explanation of action. *Philosophical Explorations*, 20(2), 242-258.
- Butterfill, S. A., & Sinigaglia, C. (2014). Intention and motor representation in purposive action. *Philosophy and Phenomenological Research*, 88, 119–145.
- Chambon, V., & Haggard, P. (2012). Sense of control depends on fluency of action selection, not motor performance. *Cognition*, 125(3), 441–451.
- Chambon, V., Sidarus, N., & Haggard, P. (2014). From action intentions to action effects: How does the sense of agency come about? *Frontiers in Human Neuroscience*, 8, 320.
- Christensen, W, Bicknell, K. McIlwain, D., & Sutton, J. (2015) The Sense of Agency and Its Role in Strategic Control for Expert Mountain Bikers. *Psychology of Consciousness: Theory, Research, and Practice*, 2, 3: 340-353.
- Christensen, W, Sutton, J. & McIlwain, D. (2016) Cognition in Skilled Action: Meshed Control and the Varieties of Skill Experience, *Mind & Language*, 31, 1: 37–66.
- Cushman, F., & Morris, A. (2015). Habitual control of goal selection in humans. *Proceedings of the National Academy of Sciences*, 112(45), 13817-13822.

- Dezfouli, A., & Balleine, B. W. (2013). Actions, action sequences and habits: evidence that goal-directed and habitual action control are hierarchically organized. *PLoS computational biology*, 9(12), e1003364.
- Dickinson, A. (1985). Actions and habits: The development of behavioural autonomy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 308, 67e78.
- Douskos, C. (2017). The spontaneousness of skill and the impulsivity of habit. *Synthese*, 1-24.
- Dreyfus, H., & Dreyfus, S. E. (1986). *Mind over machines*. New-York: Free Press.
- Elsner, B., & Hommel, B. (2001). Effect anticipation and action control. *Journal of experimental psychology: human perception and performance*, 27(1), 229-240.
- Endsley, M.R. (1995) Toward a theory of situation awareness in dynamic systems. *Human Factors* 37, 32–64.
- Endsley, M. R. (2006) Expertise and situation awareness. In K. A. Ericsson, N. Charness, P. J. Feltovich & R.R. Hoffman (eds). *The Cambridge Handbook of Expertise and Expert Performance*. Cambridge: Cambridge University Press.
- Ericsson, K. A. (2006). The influence of experience and deliberate practice on the development of superior expert performance. In K. A. Ericsson, N. Charness, P. J. Feltovich & R.R. Hoffman (eds). *The Cambridge Handbook of Expertise and Expert Performance*. (pp. 683-703). Cambridge: Cambridge University Press.
- Fitts, P. M., & Posner, M. I. (1967). *Human performance*. Belmont, CA: Brooks/Cole.
- Fridland, E. (2014). They've lost control: Reflections on skill. *Synthese*, 91(12), 2729–2750.
- Fridland, E. (2015). Automatically minded. *Synthese*. doi:10.1007/s11229-014-0617-9.
- Fridland, E. (2017). Skill and motor control: intelligence all the way down. *Philosophical Studies*, 174(6), 1539-1560.
- Fridland, E. (2019) Longer, smaller, faster, stronger: On skills and intelligence, *Philosophical Psychology*, 32:5, 759-783

Haith, A., & Krakauer, J. (2013). Theoretical models of motor control and motor learning. In A. Gollhofer, W. Taube, & J. B. Nielsen (Eds.), *Routledge handbook of motor control and motor learning* (pp. 1–28). USA: Routledge.

Haith, A. M., & Krakauer, J. W. (2018). The multiple effects of practice: skill, habit and reduced cognitive load. *Current opinion in behavioural sciences*, 20, 196-201.

Hardy, L., Barlow, M., Evans, L., Rees, T., Woodman, T., & Warr, C. (2017). Great British medalists: psychosocial biographies of super-elite and elite athletes from Olympic sports. *Progress in Brain Research*, 232: 1-119.

Jeannerod, M. (1997). *The cognitive neuroscience of action*. Oxford, UK: Blackwell.

Jeannerod, M. (2006). *Motor cognition: What actions tell the self*. Oxford: Oxford University Press.

Kahneman, D. (2003). A perspective on judgment and choice: mapping bounded rationality. *American psychologist*, 58(9), 697.

Kool, W., Gershman, S. J., & Cushman, F. A. (2017). Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychological Science*, 28, 1321e1333.

Kool, W., Cushman, F. A., & Gershman, S. J. (2018). Competition and cooperation between multiple reinforcement learning systems. In *Goal-directed decision making* (pp. 153-178). Academic Press.

Lau, H. C. & Rosenthal, D. M. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15(8): 365-373.

LeDoux, J., & Daw, N. D. (2018). Surviving threats: neural circuit and computational implications of a new taxonomy of defensive behaviour. *Nature Reviews Neuroscience*, 19(5), 269.

Levy, N. (2017). Embodied savoir-faire: knowledge-how requires motor representations. *Synthese*, 194(2), 511-530.

- Logan, G. D. (1982). On the ability to inhibit complex movements: A top-signal study of typewriting. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 778–792.
- Mazzoni, P. & Krakauer, J. W. (2006). An implicit plan overrides an explicit strategy during visuomotor adaptation. *The Journal of Neuroscience*, 26(14): 3642-3645.
- Metcalf, J., & Greene, M. J. (2007). Metacognition of agency. *Journal of Experimental Psychology: General*, 136(2), 184–199.
- Montero, B. (2016). *Thought in Action*. Oxford, UK: Oxford University Press.
- Moors, A. (2016). Automaticity: Componential, causal, and mechanistic explanations. *Annual Review of Psychology*, 67, 263-287.
- Moors, A., & De Houwer, J. (2006). Automaticity: a theoretical and conceptual analysis. *Psychological bulletin*, 132(2), 297.
- Mylopoulos, M & Pacherie, E. (2018) Intentions: The dynamic hierarchical model revisited. *WIREs Cognitive Science* 2018;e1481.
- Nadal, R. & Carlin, J. (2011). *Rafa: My story*. New York: Hyperion
- Pacherie, E. (2008). The phenomenology of action: A conceptual framework. *Cognition*, 107, 1:179-217.
- Pacherie, E. (2018). Motor intentionality. In A. Newen, L. de Bruin & S. Gallagher (eds), *The Oxford Handbook of 4e Cognition* (pp. 269-388). Oxford; Oxford University Press.
- Papineau, D. (2013). In the zone. *Royal Institute of Philosophy Supplement*, 73, 175–196.
- Papineau, D (2015). Choking and the Yips. *Phenomenology and the Cognitive Sciences*, 14: 295–308.
- Pashler, H. (1994). Dual-task interference in simple tasks: data and theory. *Psychological bulletin*, 116(2), 220.
- Pavese, C. (2015) Practical senses. *Philosophers' Imprint*, 15, 29.
- Pisella, L., Grea, H., Tilikete, C., Vighetto, A., Desmurget, M., Rode, G., Boisson, D. & Rossetti, Y. (2000). An ‘automatic pilot’ for the hand in human posterior parietal cortex: toward reinterpreting optic ataxia. *Nature neuroscience*, 3(7), 729.

- Posner, M.I. & Snyder, C.R.R. (1975). Attention and cognitive control. In R. L. Solso (ed.) *Information processing and cognition: the Loyola Symposium* (pp. 55-85). Hillsdale, NJ: Erlbaum.
- Prinz, W. (1997). Perception and action planning. *European Journal of Cognitive Psychology*, 9, 129–154.
- Ryle, G. (1949/2009). *The concept of mind*. Chicago: The University of Chicago Press; reprinted as 60th anniversary edition, London: Routledge, 2009.
- Schack, T. (2004). The cognitive architecture of complex movement. *International Journal of Sport and Exercise Psychology*, 2, 403–438.
- Schmidt, R.A. (1975). A schema theory of discrete motor skill learning. *Psychological Review* 82(4): 225.
- Schmidt, R.A. (2003). Motor schema theory after 27 years: reflections and implications for a new theory. *Research Quarterly for Exercise and Sport* 74(4): 366–375.
- Schmidt, R., & Lee, T. (2014). *Motor Learning and performance, 5th edition. E with web study guide: from principles to application*. Champaign, IL: Human Kinetics.
- Shepherd, J. (2019). Skilled action and the double life of intention. *Philosophy and phenomenological research*, 98(2), 286-305.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological review*, 84(2), 127.
- Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3.
- Stanley, J. (2011). *Know how*. Oxford: Oxford University Press.
- Stanley, J., & Krakauer, J. W. (2013). Motor skill depends on knowledge of facts. *Frontiers in Human Neuroscience*, 7(503), 1–11.
- Stanley, J., & Williamson, T. (2001). Knowing How. *Journal of Philosophy*, 98: 411–44.

Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. New York: The Macmillan Company.

Watson, P., & de Wit, S. (2018). Current limits of experimental research into habits and future directions. *Current Opinion in Behavioural Sciences*, 20, 33-39.

Wenke, D., Fleming, S. M., & Haggard, P. (2010). Subliminal priming of actions influences sense of control over effects of action. *Cognition*, 115(1), 26–38.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science* 29, 1880–1882.

Wolpert, D. M., & Kawato, M. (1998). Multiple paired forward and inverse models for motor control. *Neural networks*, 11(7-8), 1317-1329.