



Introduction: Habitual Action, Automaticity, and Control

Juan Pablo Bermúdez^{1,2} · Flavia Felletti^{3,4}

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

William James once said that “ninety-nine hundredths or, possibly, nine hundred and ninety-nine thousandths of our activity is purely automatic and habitual, from our rising in the morning to our lying down each night” (1899, pp. 65–66). Even if his estimation was wrong by half, habitual action would still be a tremendously pervasive feature of our agency. And yet, references to habitual action have been marginal at best in contemporary philosophy of action.

This neglect is due, at least, to the combination of two ideas. The first is a widespread view of habit as entirely automatic, inflexible, and irresponsive to reasons. The second is philosophy of action’s tendency (dominant at least since Anscombe and Davidson) to focus on explaining action by reference to reasons. Arguably, if habitual behavior is reasons-irresponsive, and if action is explained by reference to reasons, the study of habit would have very little to teach about action.

Recently, however, there has been a surge in philosophical interest on habit and habitual action. Novel approaches are challenging the two ideas mentioned above, arguing that (1) habitual behavior is not entirely automatic or inflexible, but instead has a particular kind of flexibility and intentionality; and that (2) acting out of habit can count as a form of acting for reasons, even in the absence of the traditional rationalizing mental states: belief, desire, and intention.

The essays contained in this issue move discussions forward in exciting new directions. In what follows we present each paper and situate it within its broader theoretical

context, so that this introduction may serve also as an introduction to the topic of habitual action. A crucial lesson that emerges from these essays is a need to move past disputes between philosophical schools or traditions. Rather than combatting between different philosophical factions, an ecumenical approach capable of skillfully bringing together elements from different traditions seems better able to tackle problems that remain unsolved. These problems include (but are not limited to): developing an account of responsibility for habitual action; explaining our ability to perform joint habitual actions; and clarifying the link between habit and self-control.

In what follows we discuss (Sect. 1) diverse approaches to the intentionality of habitual action; (Sect. 2) the relationship between habit and skill; (Sect. 3) the possibility that habits can rationalize action; and (Sect. 4) open questions and new directions.

1 Habitual Action, Automaticity, and Intentionality

As an initial definition, habits are dispositions to produce a certain well-practiced behavior in the presence of familiar cues associated with that behavior (Wood & Rünger 2016). It may seem natural, given this view of habit, to understand habitual action as a rather mindless and automatic reaction to a familiar stimulus. A kind of acquired reflex, habitual action requires no prior deliberation, no goal representation, and no attention to what one is doing. Such intuitive understanding is reflected in views like Ryle’s, according to which habits are “single-track dispositions”, generated through mindless drill and rote repetition, which produce uniform manifestations. From this perspective, an automatic action is a “mere replica of predecessors”, an automatic copy of past automatic behaviors (Ryle 1949/2009, ch. 2).

But more care is needed here. The term ‘automatic’ is associated with many different properties (unintentional, uncontrolled, inflexible, goal-independent, stimulus driven, etc.), and each one of those properties comes in degrees and

✉ Juan Pablo Bermúdez
juanpa@gmail.com

¹ Institut de Philosophie, Université de Neuchâtel. Espace Tilo-Frey 1, 2000 Neuchâtel, Switzerland

² Programa de Filosofía, Universidad Externado de Colombia, Bogotá, Colombia

³ Department of Philosophy, University of Duisburg-Essen, Essen, Germany

⁴ Department of Psychology, Leuphana University, Lüneburg, Germany

can be manifested independently of the others (Moors & De Houwer 2006). The graded character of automaticity allows us to conceive of habitual actions as displaying different levels of automaticity, or displaying some forms of automaticity without displaying others. Habit may always involve some degree or aspect of automaticity (acquired through the history of action repetition), but not all habitual actions need to be entirely automatic, mindless, and inflexible.

In fact, saying that habitual actions are mere replicas of past actions is not accurate. We have picked coffee cups many times before, but (almost) every time we do it our behavior seamlessly adapts to the current circumstances, even when doing so entirely out of habit. We need not think about the cup's precise position, the strength of our grip or the arm's movement from table to mouth—these issues are sorted out automatically, but every time slightly differently, as the habitual action adapts to each particular circumstance. In short, habitual actions display flexibility within the narrow bounds of familiar contexts.

Thus some space opens to conceive of habitual actions as endowed with some kind of intentionality: they display the agent's ability to behave appropriately in a narrow range of varying but familiar contexts, often in goal-directed ways, and revealing some sort of learning acquired through experience. But how can we account for said intentionality, especially if it is expressed in largely automatic behavior?

One way to characterize habitual action's intentionality is to focus on habits themselves, in isolation from representational mental states and the processes that involve them. The structure and functioning of habit—so the idea goes—is in itself sufficient to guide the agent as she copes intelligently with her environment. Call that *non-representationalism*. Alternatively, one could try to account for the intentionality of habitual action by appealing to mental representations different from the traditional beliefs and desires, which can guide non-deliberative, habitual action. Call this *representationalism*. In what follows we introduce different non-representationalist and representationalist approaches, and later discuss whether the contrast between them is really as drastic as it sounds.

1.1 Non-representational Approaches

While non-representationalism is not new, it has been brought back into the contemporary debate by Bill Pollard (2006), who argues that habitual actions have a certain *intrinsic intentionality*, a kind of teleology acquired through repetition that they seek to fulfill. While this idea may sound promising, Pollard offers no clear explanation of what underlies that intrinsic teleology. However, the view can be fleshed out by referring back to earlier traditions that have developed versions of it in detail.

One of these is the pragmatist tradition, and particularly John Dewey's work, which Miyahara and Robertson (this volume) describe and update. According to them, Dewey holds that intentionality is a property not of an individual habit, but of *networks of habits* working in tandem with well-structured practical environments. The environment and the agent's network of habits form a sort of rich symbiosis capable of structuring complex intelligent behavior. Miyahara and Robertson discuss the example of driving home along the habitual route guided by the signs of a well-structured road: agents can accomplish this thanks to a set of habits cultivated through their driving history. As long as the habits are in sync with an environment of traffic signs and well-delineated lanes, agents need not use any explicit rules or goal representations to intelligently navigate the space.

Importantly, this proposal allows us to distinguish between intelligent habits, which enable intentional agency, and unintelligent habits, which do not. The latter lack the richness of connections with a network of other habits and with a well-structured environment, and thus tend to produce rote repetition and inflexible behavior. A possible analysis of tics or addiction could stem from this distinction.

In a similar vein, the phenomenological tradition also provides rich elements to flesh out the claim that habits themselves are endowed with intentionality. Merleau-Ponty's notion of *motor intentionality* is particularly relevant. In *Phenomenology of Perception* Merleau-Ponty (1945/2012) discusses the case of Schneider, a German World War I ex-combatant who, after receiving serious brain injuries, displayed an intact ability to perform habitual actions but severely compromised capacities to perform 'abstract movements' [arbitrary, situation-independent movements, like pointing to one's nose (as opposed to grabbing one's nose)]. From this case Merleau-Ponty draws the inference that *motor intentionality* (the intentionality at play in our routine, unreflective bodily actions) is independent from *cognitive intentionality* (the representation-dependent, reflective control of action).

Building on these ideas, Cappuccio et al. (this volume) propose that motor intentionality produces intelligent behavior whose success conditions are specified, not by goal representations, but by the habitual action's properly responding to the environment's affordances, which solicit the agent to act in a certain way. Such intelligent responsiveness is enabled by "sensorimotor habits" (well-trained dispositions to produce familiar solution-types to known situation-types), whose flexibility is based not on representations, but on the agent's *body schema* (another Merleau-Ponty term): a system of transposition of motor tasks that finds equivalent positions in different orientations.

These pragmatism- and phenomenology-inspired views propose interesting ways of fleshing out the idea that habitual actions display an intrinsic, representation-independent

teleology. But some challenges exist. Consider, first, that explaining the intelligence of habitual actions involves not only explaining their flexibility (e.g. my ability to grab differently sized cups, or to type my password with one hand or two hands), but also their sensitivity to, and ability to coordinate with, higher-order cognitive states like intentions (Ferretti & Zipoli Caiani, this volume). In fact, many habitual action performances are sensitive to our occurrent goal representations: we can e.g. write our signature more slowly or more quickly, depending on how important it is to get it right; we can recruit our driving habits to take ourselves home or to a restaurant we have never been to. Non-representational approaches face the challenge of explaining how this integration between motor and cognitive intentionality works (Pacherie 2018), and the claim that one is non-representational and the other representation-based makes this integration harder to explain.

Additionally, providing an account of habitual actions arguably requires also explaining how these actions make sense *from the agent's point of view* at the time of acting (Kalis & Ometto, this volume), and explaining in which sense an action belongs to her agent, rather than to someone else or to elements of her environment (Asma, this volume). This might require referring to the agent's goals or reasons, or at least explaining how the action coheres with or derives from them. The issue becomes practically important when it comes to holding an agent responsible for her habitual action. It is unclear whether an explanation that appeals merely to the links between environmental cues and networks of habits, or to environmental affordances soliciting responses from the agent, can succeed in rendering the action intelligent and intelligible from the agent's point of view, and therefore help in grounding responsibility attributions, without any reference to the agent's goals, values, and plans.

1.2 Representational Approaches

An alternative path is explaining the intentionality of habit by positing representation types particular to non-deliberative, habitual action. This approach holds the promise of being able to account for the integration between intentions and habitual behavior, by specifying how the representations constituting intentions get articulated with representations that produce and control habitual action. That, however, has proven to be a hard challenge, since it requires solving what has been called the *interface problem* (Butterfill & Sinigaglia 2014): assuming that habitual behavior is guided by motor representations developed through repeated practice, how do those motor representations (which are fine-grained, non-propositional, and motorically formatted) interlock with beliefs, intentions, and the like (which are abstract and propositional)?

While there is a burgeoning literature on this issue (for a review see Mylopoulos & Pacherie 2019), not much has been written specifically with habitual action in mind—until now. Ferretti and Zipoli Caiani (this volume) propose that the solution lies in positing *action concepts*: representations which develop through repeated practice and, being conceptual, can form part of the content of intentions, but also have a motoric format, which makes them able to trigger the motor commands necessary to execute and control habitual actions.

Pacherie and Mylopoulos (this volume) offer another proposal. Simple motor actions are represented by *motor schemas*: representations of the action's invariant features, including its functional organization, relative timing, and spatial configuration. A motor action's non-structural aspects correspond to the motor schema's parameters, which can be adapted to fit the specific context. Through repeated practice, agents chunk multiple motor schemas together into what Pacherie and Mylopoulos call *structured action representations* [SARs]: hierarchically organized assemblages of motor schemas that allow the agent to produce and control more complex motor action units. The bridge between intentions and motor representations is possible because the agent can exert intentional control over the parameters of the motor schemas that compose the SAR. This kind of architecture is intended to explain expert skillful action, but can also be used to distinguish between skills and habits. We will come back to this point.

This representational approach to habitual action has many attractive aspects. For one, it offers an explanation for how habitual actions simultaneously mix automaticity and flexibility. 'Action-oriented' representations (like Ferretti and Zipoli Caiani's action concepts, or Pacherie and Mylopoulos' SARs) allow for the automatization of complex routines while leaving a reduced number of parameters still susceptible to intentional control. Additionally, representational approaches also promise to shed light on how the intelligence of habit can combine with that of higher-order representations and deliberation. But in this regard the work is far from done. The interface problem remains a matter of great debate, and there is much disagreement about the specific nature of the 'intermediary' representations (which integrate higher-order intentions and lower-order motor representations) and of the integration process itself.

Another open question: some researchers think that habitual actions are always triggered by an intention. According to this view, a habitual action *H* would be a way of *F*-ing, where you can fulfill your intention to *F* by *H*-ing (Amaya 2020). But while that may happen very often, habitual actions can also be produced without the contribution of intentions, when situational cues activate

action-oriented representations without the mediation of agent-avowed intentions. The intentionality of intention-implementing habitual actions is easy to explain, since it ‘flows downstream’ from the intention; but how do representational approaches explain the intentionality of intention-independent habitual actions?¹

This dovetails with a possible complaint that representational approaches fail to account for the role of the environment in guiding and structuring habitual actions. Since these are behaviors triggered in response to familiar environmental cues, said cues should probably play a role in explaining why some habitual actions make sense and others do not, like Miyahara and Robertson (this volume) argue. But representationalist approaches have so far not much to say about the context’s contribution.

In the end, we think a combination between representationalist and non-representationalist positions is not only possible, but desirable. After all, they both agree that the intentionality of habitual action cannot be explained by reference only to propositional and conceptual representations. The representations posited by the representationalist positions discussed here are non-conceptual or non-propositionally formatted, and they are tailor-made for guiding, monitoring, and controlling motor processes. One may not need to reject the latter kind of representation to capture the embodied and situated nature of motor intentionality. (Is there that much of a difference between Pacherie and Mylopoulos’ *motor schemas* and Merleau-Ponty’s *body schema*?) Perhaps an ecumenical approach can give important explanatory roles to both non-propositional, motor representations and to the structures and affordances in the practical environment which contribute to non-deliberately guiding the agent’s behavior. This seems to us to be the most promising way forward.

2 Interlude: Habit vs. Skill

There is a significant difference between the dispositions and abilities underlying the habitual nail biter’s nervous nibbling and the world-class athlete’s gold-medal performance. Yet many authors simply make no distinction between habit and skill. This is perhaps in part because the most popular dividing line, which goes back to Ryle, seems highly problematic. According to Ryle (1949/2009), skills are multi-track dispositions (flexible, “intelligent capacities” that constitute exercises of ‘knowing-how’), whereas habits are

¹ While action-oriented representations can be used to guide action, they do not count here as intentions. Intentions are understood as propositionally formatted and having conceptual content, whereas action-oriented representations are understood as being either loose concepts (e.g. Ferretti and Zipoli Caiani’s action concepts) or non-propositionally formatted and having non-conceptual content (e.g. Pacherie and Mylopoulos’ SARs).

single-track dispositions acquired through mindless drill. Whereas acting out of habit implies acting automatically and without vigilance or care to what one is doing, acting skillfully implies being mindful, attentive to errors and aiming toward improvement.

This radical opposition does not work as it stands, given several previously discussed points: the scalar nature of automaticity and attention (a behavior can be more or less automatic, performed with more or less attention); and the fact that habitual actions can also display flexibility (like my typing my password with one or two hands). Add to this that skilled agents can perform skillful actions with vigilance or without (Martens, this volume).

A promising alternative distinction was recently proposed by Douskos (2019): although habit and skill both involve automaticity, they are automatic in two different senses: the automaticity of habit is *impulsivity*: the tendency to respond to familiar circumstance in a specific, well-practiced way; but the automaticity of skill is *spontaneity*: a capacity to respond to a range of circumstances in ways conducive to the agent’s goals. The spontaneity of skill requires attention for the agent to find the right way to respond in the given circumstance, whereas the impulsivity of habit is independent from attention. In fact, to the extent that a habitual action is impulsive, it is performed unattentively. The automaticity of skill involves attention, but the automaticity of habit repels attention.

Habitual actions can, under certain circumstances, be performed less impulsively and more attentively (I can pay close attention when I am making my coffee this morning because I want to impress my guest), but when we act habitually we often act with a large degree of impulsivity, and that entails paying little to no attention to how we are acting or how the action fits with our current goals.² Impulsivity thus exposes agents to the risk of making mistakes when acting habitually: they can slip by responding to familiar contexts in familiar ways when it is not appropriate to do so. Kalis and Ometto (this volume) mention the brilliant example of a group of Soviet soldiers on a secret mission to Cuba who, upon landing, marched away from the plane in perfect formation. We will return to this issue below, since

² Notice that ‘impulsivity’ in the sense presented here need not involve speed, rashness, or violent movement. You can act impulsively and *slowly* (e.g. I have a habit of slowly looking up at ceilings whenever visiting a church), and even *omissions* can be impulsive (e.g. I can habitually omit to turn on the light when I go to my usual bathroom). What makes a habit impulsive in the relevant sense is just that it tends to produce a familiar response in a familiar circumstance.

it has consequences for thinking about the normativity of habitual actions.³

Another line of separation between skill and habit could be found at the level of the representations and cognitive processes involved. Following Pacherie and Mylopoulos, we could say that both habit and skill require SARs, but those constitutive of skill—being developed not through mere repetition but *deliberate practice* (Ericsson et al. 2009)—are of a higher quality, thus allowing the expert agent better movement selection, effect anticipation, monitoring and control throughout a broader range of situations.

According to Pacherie and Mylopoulos (this volume), the psychology of expert agents is different also because they have better decision-making routines. When choosing how to implement a given action in a familiar context, we can decide whether to use an internal model (a representation of the context's causal structure and our agentic capacities, through which we can make predictions of action outcomes) in order to identify the optimal action patterns, or to go 'model-free' and simply implement the tried-and-tested habitual response to the familiar context. Model-free, or habitual, action selection is faster and less cognitively costly (since it does not involve producing mental predictions), but model-based action selection is more flexible (since it can better adapt behavior to the current circumstance's properties). Pacherie and Mylopoulos suggest that highly skilled agents have better decision procedures to navigate this "efficiency-flexibility tradeoff", i.e. to choose when it makes sense to go model-free and when model-based. Habits give us only a model-free solution: there is no model-based alternative. (Notice that this echoes Douskos' distinction between impulsivity and spontaneity).

3 Habitual Action and Acting for Reasons

3.1 A Challenge from Habitual Action

As Pollard has pointed out, we sometimes explain actions by referring to habits, e.g.: "He missed the turn because he carried straight out of habit", "He's biting his nails because he has that bad habit" (Pollard 2006). A crucial characteristic of habit explanations is that they make no reference to the psychological states that appear in traditional action

explanations: beliefs and desires. This leads Pollard to categorize them as non-psychological explanations.

For some, the non-psychological nature of habit explanations turns habitual action into a potential counterexample to the dominant *causal theory of action*, according to which a behavior is an action (and not a mere bodily movement) just in case its production involves the correct unfolding of certain psychological processes. States like beliefs and desires are usually involved in explanations of such processes. Influential causal theories (Davidson 1963) hold that psychological states are not only the causes of the action, but also the agent's *reasons* for performing the action: they both cause the action and explain why the agent performed it, i.e. why it made sense from her point of view.

It could seem to follow from this causalist line of reasoning that an action *F*'s being *responsive to reasons* depends on the agent's consciously and explicitly considering those psychological states that constitute her reasons, and on her choosing to *F* because of that process of consideration. Clearly, habitual action usually involves no conscious consideration of reasons to act, no explicit choice to perform the habitual action. So, if habitual actions are intentional actions, and explanations of those actions are non-psychological, they are a counterexample to causalism (Pollard 2006; Kalis & Ometto, this volume).

3.2 Replies to the Challenge

Could a causalist reply by maintaining that habits are psychological entities just like beliefs or desires? Pollard (2006) would disagree: psychological entities of the relevant kind are those that give us "first-person privilege". Agents seem to have privileged access to their beliefs and desires, but many habits are invisible to the agent, and often easier to discover from second-person observation. Additionally, the psychological entities that caused the action must be consciously accessible to the agent, since otherwise they cannot explain how the action makes sense from the agent's perspective (Kalis and Ometto, this volume).

That said, the causalist still has some moves available. Consider, for instance, Railton's (2009) view that our ability to grasp reasons extends far beyond our ability to explicitly assess beliefs and desires, and to deliberately make a choice about what to do—in other words, the view that *reasons-responsiveness* does not require *deliberation*. Railton offers cases of "fluent agency" in which we act in accordance with reasons just as competently as fluent speakers of a language use the rules of grammar. Consider the experienced driver's smooth downshifting or the home cook's meal prepping without consulting recipes. They do not explicitly think about the reasons they are following, but are still *sensitive* to them, and can respond to them without explicit reflection.

³ This way of distinguishing habit and skill suggests that agents cannot slip while performing skillful actions. Rather, the dependence of skill on a particular kind of attention makes agents susceptible to different kinds of agency breakdowns, like choking and yipping. For discussion see Bermúdez (2017).

Railton argues that our fluent sensitivity to reasons is enabled by fast, automatic evaluations of the situation that have been developed through habituation and socialization processes since early age. These automatic processes may themselves be inaccessible to the agent, but they generate affective experiences that the agent *does* have conscious experience of. (I am about to tell a story in conversation, but suddenly sense that the topic is making my interlocutors uncomfortable, so I stop myself). We develop these affective, context-sensitive evaluations via habit acquisition, and particularly our processes of socialization and enculturation (Hufendiek 2020). Hufendiek calls these affective reactions “action-oriented representations”, since they prepare the agent to deal with the situation at hand in a reasons-responsive way. Such reasons-responsiveness comes from the cumulative learning processes that have sedimented into the agent’s habits throughout her development.

A causalist account of reasons-responsiveness that relies on action-oriented representations offers one promising line of response to the challenge from habitual action. First, the fact that habit explanations do not mention beliefs or desires does not entail they are non-psychological. The reference to habits may implicitly include a reference to habit-mediated psychological states like action-oriented representations. The causalist can thus still say that we are able to generate reasons-responsive habitual actions without deliberation, because psychological processes involving such accessible feelings and intuitions produce reasons-sensitive intentional actions.

There is yet another path available to the causalist. As mentioned above, we often perform a habitual action as a way to implement an occurrent intention: we drive home by taking the usual route or eat cereal because it is our habitual breakfast. Thus, habitual actions can be means-actions through which we implement an intention. According to Amaya (2020), these cases show that actions can be responsive to reasons thanks to habits *in absence of beliefs*. Given an intention to *F*, I can guide the intention’s implementation through my belief that *G*-ing is a way of *F*-ing (in which case that belief rationalizes my *G*-ing); or I can just follow the habit linking my intention to *F* with its default implementation *H* (in which case my habit rationalizes my *H*-ing). In the latter case, my habit rationalizes my action by linking an intention and its default, well-practiced implementation. If the habit was developed via a history of successful reinforcement, this history can serve as inductive support for the habitual action’s rationality (Amaya 2020).

There are, of course, cases in which habits mistakenly guide our actions, like when the driver misses the turn because she continued straight out of habit. These slips demonstrate that habits can guide action independently

from beliefs, since slips, according to Amaya, are intentional actions that are contrary to an occurrent belief about what one is doing. (She believes she is driving to her friend’s house, but the habit is guiding her back home).

These are only two possible causalist replies to the challenge from habitual action. Their merits must be critically assessed, but their availability suggests that we should reinterpret the target of the habitual action challenge. Initially we saw it as a challenge to causalism in general; but the existence of intentional habitual actions seems to threaten *particular forms* of causalism, i.e. forms that hold either of these views:

- *Deliberative guidance*: An action *F* is reasons-responsive if it is caused and explained by an explicit deliberation about the reasons to *F*.
- *Doxastic guidance*: In order to implement an intention in a reasons-responsive way, the implementation must be guided by a belief about ways to implement it.

The challenge from habitual action suggests that causalism must abandon *deliberative guidance* and *doxastic guidance*, and explain how we are able to act for reasons in the absence of explicit deliberation. But neither of those claims is a necessary component of causalism.

3.3 Other Approaches

Non-causalist accounts of how habits rationalize actions are also available, and provide further insights into the rationality of habit. In his seminal work, Pollard (2006) proposed that habit explanations are *constitutive explanations*, in two senses:

- (1) *Habitual actions are constitutive of habit*. Each habitual action is a part of the habit, and the habit is constituted by its manifestations in particular habitual actions.
- (2) *Habits are constitutive of the agent*. Echoing Dewey’s claim that habits “constitute the self” (Dewey 1922, p. 25), Pollard claims that having a certain habit is part of what it is to be the specific agent one is. Especially for habits we cultivate long-term, once we have acquired them, they help constitute our identity.

The second constitutive relationship is developed and sharpened by Wagner (this volume), who offers an account of “identity-shaping habits”: habits that are shaped through long-standing, repeated practice, and which thereby acquire a role in the story we tell ourselves about who we are. Given their role in our self-narrative, we have a standing interest in managing their expressions (by either up- or down-regulating the habitual actions that manifest them). Wagner

puts the notion of identity-shaping habits to work in tackling philosophical problems surrounding personal identity.

Given the two constitutive links, a habit explanation can be said to rationalize habitual action by situating it as part of a broader pattern of the agent's behavior, and by presenting that pattern as constitutive of the agent's identity. That said, a critic might raise the challenge that the habit-agent link holds only in special cases, like Wagner's identity-shaping habits, but cannot apply to more mundane habits less central to the agent's self-conception. A nail-biting habit or the tendency to put on the left sock before the right in the mornings don't play much of a role in defining our identity. The rationalization of these more mundane habitual actions can thus rely only on the first constitutive relationship (between action and habit); but this might compromise the link between the mundane habitual action and the agent's self-identity. And if there is no link to self-identity the constitutive explanation may fail to make sense of the habitual action from the agent's point of view.

Asma (this volume) raises an additional challenge for Pollard's constitutivism, arguing that it cannot distinguish between habitual actions and mere habitual bodily movements (i.e. abitudinal movements that do not amount to action). As a test case, Asma depicts an agent who has the habit of going to physiotherapy and having her arm moved up and down by her therapist. The agent always retains "veto control" (the capacity to stop the movement at will), and the movement itself is part of the agent's physiotherapy habit. Given this, Asma holds Pollard's constitutivism would categorize that bodily movement as an action of the patient's. But it is clearly not.

Asma's explanation of the case is that there is an intrinsic difference between intentional habitual actions and mere habitual movements: only for the former the agent has *practical knowledge* of what she is doing as she does it. Here she echoes Anscombe's (1957/2000) view, which both Asma and Kalis & Ometto (this volume) apply to the case of habitual action. According to them, an intentional action makes sense from the agent's perspective because the agent herself knows *that* she is performing the action and *why* she is performing it, without needing to infer that knowledge from other, more immediate facts. What distinguishes the physiotherapist patient's movement from a habitual action is that in that case the agent has no practical knowledge of that movement. This non-observational, first-person knowledge of the agent's action is simultaneous with its performance. The existence of intentional habitual actions shows that practical knowledge is independent from prior deliberation, since they are non-deliberative and yet intentional.

Critical readers might ask whether this Anscombean approach can allow for the existence of irrational habitual actions. Take the case of the driver who missed the turn because he continued straight out of habit, or Kalis & Ometto's case of the Soviet undercover soldiers walking in formation out of habit. Are those slips intentional actions? From

an Anscombean perspective they seem not to be, since the agents lack practical knowledge of what they are doing *while* they are doing it. (The driver believes he is going to a certain location, while he is going to another; the soldiers believe they are concealing their identity while they are forming a military figure together). Thus habitual action slips fail the practical knowledge test. But we have independent reasons to think slips are intentional actions—they are mistakes, but still they are intentional, since they are guided by a habit. A satisfactory account of habitual actions must have room for irrational, yet intentional, habitual actions; whether a pure Anscombean approach can do this is an open question.⁴

To conclude this section, the discussion of how we can explain habitual actions as acting for reasons has served to reveal the limitations of certain causalist approaches to action, specifically those committed to *deliberative guidance* or *doxastic guidance*. But in this respect (i.e. in the rejection of *deliberative guidance* and *doxastic guidance*), current causalist, Anscombean, and constitutive approaches turn out to be closer to each other than to traditional causalism. It seems then that, more than inviting us to reject causalism *tout court*, habitual actions invite us to reconsider the role of deliberation and belief in action guidance.

Each of the approaches presented here offers unique contributions and faces its own challenges. So a lesson to be drawn here again is that we can be more ecumenical about our theoretical orientations (whether one's view is causalist or Anscombean can take the backseat) and draw from multiple traditions to tackle problems that are still in need of solution and common to multiple theoretical orientations.

4 New Directions

4.1 Habitual Weakness

A particular type of habitual action slips is what Silver (2019) calls "habitual weakness". These are cases in which an agent acts against her better judgment or her resolution, not because of a strong wayward motivation (like in traditional cases of akrasia or weakness of will), but because of a habit. E.g. I realize I am running late for my appointment, so I commit to going about my morning routine more quickly than usual; but then I catch myself taking just as long as always in the shower.

⁴ One might reply that a slip is intentional under one description (driving the car) but not under another (missing the turn). The problem, as Amaya (2013) points out, is that slips are intentional under at least one description under which the agent would have preferred not to have performed it (driving home along the habitual route). What makes the action intentional under that non-preferred description is that the agent allowed her habit to guide the action.

Silver argues that actions displaying habitual weakness are cases of akrasia and weakness of will. If he is right about that, it would entail that emotions and motivations are not necessary components of akrasia—habits would be sufficient to trigger it. The possibility that habitual action can constitute akratic behavior in the absence of motivation would require a radical rethinking of the nature of akrasia, weakness of will, and self-control.

4.2 Responsibility for Habitual Action

The fact that habitual actions can misfire, as in cases of habitual weakness and slips, highlights the issue of responsibility. If I commit to driving my neighbor to the airport but on the way I slip and take the route to my office, upon which she loses her flight, she could reasonably hold me responsible. Douskos (this volume) argues there is a gap in the responsibility literature in this regard. He assesses whether responsibility for habitual actions can be explained by some forms of *control theory*, which holds that a morally competent agent is responsible for *F*-ing (where *F* can be an act or omission) when the agent's control played a role in bringing *F* about, and she is aware of *F*'s moral significance while *F*-ing. Habitual actions raise a challenge to control theories because of the latter condition, which seems to demand awareness from the agent, since many habitual actions are performed without attention.

Some varieties of control theory are promising since they have been tailored to account for *unwitting omissions*, i.e. unintentional omissions of which the agent is unaware at the time they are occurring. It seems clear we can be responsible for unwitting omissions (I accept to water my neighbor's plants while she is away, but I forget. If the plants die, my neighbor would be right to hold me accountable). Douskos notices that unwitting omissions and habitual actions are problematic for the same reason: in both cases the agent is unaware (because her attention is elsewhere) of her conduct's morally significant features. He assesses whether two versions of control theory meant to explain responsibility for unwitting omissions (tracing and capacitarian accounts) can also explain our responsibility for habitual actions, and concludes that neither can successfully do so. If he is right, then a theory of responsibility for habitual actions remains sorely needed. Considering the pervasiveness of habitual action in everyday life, this is a big gap in need of filling (see Murray & Vargas 2020, for a possible starting point).

4.3 Acting Habitually Together

When we act habitually, we do not always act alone. Many habitual actions and routines are shared, and we engage in them together. Martens (this volume) offers a way to

account for our capacity to perform joint habitual actions. She notices that available accounts of joint action start from the assumption that the agents have a shared goal, and then explain what allows them to pursue that goal together. In contrast, Martens develops an account of coordination that does not depend on a prior agreement about a shared goal. She proposes an "interaction-dominant account", which allows for habitual joint actions to emerge bottom-up, in the absence of a shared goal representation.

Martens discusses cases of actions that the agent does not initiate, but picks up from the cues in the behavior of others and in the material situation, and then carries forward in her own behavior. (E.g., a new police officer notices cues suggesting that her more experienced colleagues treat white people in a certain way and non-white people in another, and picks up these behaviors without setting an explicit goal to do so). Such "intentional continuings", when they become engrained in a shared routine, constitute cases of joint habitual actions that emerge bottom-up, without agents sharing a goal to be pursued together. Martens' analysis offers new conceptual tools for understanding how phenomena like institutional routines and group identities might originate. These concepts would also help explain why routines can be so "sticky"—once they come about, they can resist reform since they are goal-independent—, and would ground the possibility, mentioned by Silver (2019), that habitual weakness can be displayed also at the group level.

5 Closing Remarks

Habitual action is an 'amphibious' phenomenon. It is automatic behavior that is nevertheless intentional. It expresses the agent's identity but is triggered by the environment. It can cohere with our goals, but can also be insensitive to them. It is intelligent, but it is also mindless. Since habitual action straddles many traditional dichotomies (automatic and controlled, internal and external, intelligent and mindless), work on the topic promises to move us past dichotomic thinking and into a more nuanced view of action and the mind. Along the way, it has the power to show that phenomena at one point held to be central to intelligent action (like propositional representations, deliberation, and belief) are not as central as once thought.

Advancing in understanding the topic requires a view that is not very concerned with which schools of thought we endorse (whether we are representationalists or anti-representationalists, causalists or Anscombeans), but rather uses conceptual tools from multiple sources to tackle the complex problems that remain unsolved. We see the contributions to this issue as moving the conversation in that direction.

References

- Amaya S (2013) Slips. *Noûs* 47(3):559–576
- Amaya S (2020) Out of habit. *Synthese*. <https://doi.org/10.1007/s11229-020-02780-3>
- Anscombe GEM (1957/2000) *Intention*. Harvard University Press, Cambridge (**Original work published 1957**)
- Bermúdez JP (2017) Do we reflect while performing skillful actions? Automaticity, control, and the perils of distraction. *Philos Psychol* 30(7):896–924
- Butterfill SA, Sinigaglia C (2014) Intention and motor representation in purposive action. *Philos Phenomenol Res* 88(1):119–145
- Davidson D (1963) Actions, reasons, and causes. *J Philos* 40(23):685–700
- Dewey J (1922) *Human nature and conduct: an introduction to social psychology*. Henry Holt & Co, New York
- Douskos C (2019) The spontaneousness of skill and the impulsivity of habit. *Synthese* 196(10):4305–4328. <https://doi.org/10.1007/s11229-017-1658-7>
- Ericsson KA, Nandagopal K, Roring RW (2009) Toward a science of exceptional achievement: attaining superior performance through deliberate practice. *Ann N Y Acad Sci* 1172:199–217. <https://doi.org/10.1196/annals.1393.001>
- Hufendiek R (2020) Emotions, habits, and skills: action-oriented bodily responses and social affordances. In: Caruana F, Testa I (eds) *Habits: pragmatist approaches from cognitive neurosciences to social sciences*. Cambridge University Press, Cambridge, pp 100–119
- James W (1899) The laws of habit. In: James W (ed) *Talks to teachers on psychology and to students on some of life's ideals*. Henry Holt & Co., New York, pp 64–78
- Merleau-Ponty M (1945/2012) *Phenomenology of perception*. Routledge, London (**D. A. Landes, Trans Original work published 1945**)
- Moors A, De Houwer J (2006) Automaticity: a theoretical and conceptual analysis. *Psychol Bull* 132(2):297–326. <https://doi.org/10.1037/0033-2909.132.2.297>
- Murray S, Vargas M (2020) Vigilance and control. *Philos Stud* 177(3):825–843
- Mylopoulos M, Pacherie E (2019) Intentions: The dynamic hierarchical model revisited. *Wiley Interdiscip Rev: Cogn Sci* 10(2):e1481
- Pacherie E (2018) Motor intentionality. In: Newen A, De Bruin L, Gallagher S (eds) *The oxford handbook of 4E cognition*. Oxford University Press, Oxford, pp 368–388. <https://doi.org/10.1093/oxfordhb/9780198735410.013.19>
- Pollard B (2006) Explaining actions with habits. *Am Philos Q* 43(1):57–69
- Railton P (2009) Practical competence and fluent agency. In: Sobel D, Wall S (eds) *Reasons for action*. Oxford University Press, Oxford, pp 81–115
- Ryle G (1949/2009) *The concept of mind*. Routledge, London (**Original work published 1949**)
- Silver K (2019) Habitual weakness. *Thought: A J Philos* 8(4):270–277. <https://doi.org/10.1002/tht3.431>
- Wood W, Rünger D (2016) Psychology of habit. *Annu Rev Psychol* 67(1):289–314. <https://doi.org/10.1146/annurev-psych-122414-033417>

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