Review of Philosophy and Psychology
https://doi.org/10.1007/s13164-019-00452-x

Intention at the Interface

Ellen Fridland[1]

Published online: 05 December 2019

Abstract
I identify and characterize the kind of personal-level control-structure that is most relevant for skilled action control, namely, what I call, "practical intention". I differentiate between practical intentions and general intentions not in terms of their function or timing but in terms of their content. I also highlight a distinction between practical intentions and other control mechanisms that are required to explain skilled action. I'll maintain that all intentions, general and practical, have the function specifying (and thus guiding according to those specifications), sustaining, and structuring action but that several functions that have been attributed to proximal intentions are actually implemented by other control mechanisms that are not themselves best identified as intentions. Specifically, I will claim that practical intentions do not initiate, monitor, specify or guide the fine-grained, online, kinematic aspects of action. Finally, I suggest that the way in which practical and general intentions should be differentiated is in terms of their content, where general intentions specify the overall goal, outcome, or end of an action as it is conceived of by the agent at a time, and practical intentions determine the means to that end. I conclude by providing empirical evidence to support this way of characterizing the intentions that "interface" with the mechanisms of motor control. Though this discussion has repercussions for action in general, I will limit my discussion to cases of skill.

1 I. Introduction

In a 2014 paper, *Intention and Motor Representation in Purposive Action*, Stephen.

Butterfill and Corrado Sinigaglia (B&S, from hereon) introduce and propose a solution to what they call "the interface problem". The interface problem concerns the implementation of intentional action, which purportedly requires coordination between states coded in different formats. Specifically, B&S argue that both intentions and motor representations specify action outcomes. However, since these outcomes are coded in different

✉ Ellen Fridland
Ellen.fridland@kcl.ac.uk

[1] King's College London, Strand WC2R 2LS, UK

formats, i.e., propositionally in the case of intentions and in motor code in the case of motor representations, there is an open question about how those outcomes come to match or "interface" nonaccidentally. B&S claim that any "content based" solution to this interface cannot work since it would require a mysterious translation process that no one knows anything about. As such, they propose that the connection between the two kinds of states is achieved through a process akin to demonstrative reference.

Though I disagree that there is an interface problem so characterized, I do think that it is imperative to specify exactly what it is that different kinds of representational states coded in different formats contribute to purposive action, specifically, to skilled action control. So, though I myself favour a content-respecting solution to the problem of skilled action control, this solution still requires specification of the kind of states that need to communicate with one another in order to make skilled action possible.[1] And doing this requires articulating both what it is that motor representations represent and specifying how best to characterise the personal-level, intentional states that are involved in action implementation.

Accordingly, this paper will attempt to identify and characterize the kind of personal-level control-structures or intentions are most relevant for skilled action control, namely, what I will call, "practical intentions". Importantly, I will differentiate between practical intentions and general intentions not in terms of their function or timing but in terms of their content. Moreover, I will highlight a distinction between practical intentions and other control mechanisms that are required to explain skilled action. I'll maintain all intentions, general and practical, have the function specifying (and thus guiding according to those specifications), sustaining, and structuring action but that several functions that have been attributed to proximal intentions are actually implemented by other control mechanisms that are not themselves best identified as intentions. Specifically, I will claim that practical intentions do not initiate, monitor, specify or guide the fine-grained, online, kinematic aspects of action. Finally, I'll suggest that the way in which practical and general intentions should be differentiated is in terms of their content, where general intentions specify the overall goal, outcome, or end of an action as it is conceived of by the agent at time t, and practical intentions determine the means to that end. I'll conclude by providing empirical evidence to support this way of characterizing the intentions that "interface" with the mechanisms of motor control. Though this discussion has repercussions for action in general, I will limit my discussion to cases of skill.

2 Skill and Intentional Actions

In discussions of know how or skill, one paradigmatic feature of skilled actions that most philosophers agree is characteristic of skill is their intentionality or goal-directedness (Ryle 1949; Stanley and Williamson 2001; Noë 2005; Bengson and Moffett 2011; Stanley and Krakauer 2013; Fridland 2014; Pavese 2018). That is, most philosophers agree that when a skilled action is instantiated, that action has some

[1] The question of how personal-level goal states and motor control states communicate, on the account I favor, differs from the formulation of the interface problem in that I deny that (1) there must be two states represented in different codes that match in their content (2) I also disagree that translation is a mysterious process, as B&S insist. In this way, I am persuaded by arguments that Shepherd (2017a, b) has forwarded about how representational pluralism, very generally, requires translation.

essential and necessary connection to what the agent aims or intends to do. For instance, Jason Stanley and John Krakauer (Stanley and Krakauer 2013) write that, "[P]aradigmatically, skills manifest in intentional actions" and Alva Noë, though disagreeing with the basic tenets of intellectualism, which Stanley and Krakauer defend, argues that digestion cannot be an instance of know how because it isn't the right sort of thing to qualify as know how, that is, it isn't an intentional action. He writes, "I would agree that digesting food is not the sort of action that one knows how to do. But that is because digesting is not the sort of thing one does (intentionally or otherwise). Hannah doesn't digest food; her digestive system does (in her or for her). Hannah may have excellent digestion, but she is not, in that case, excellent at digesting. Digestion is not an action that a person or animal can perform; it is a process that takes place inside a person or animal" (2005, 279). Thus, we can infer that Noë is committed to the view that instances of know how must be actions that an agent performs—actions that are intentional, goal-directed, or done on purpose.

And even Gilbert Ryle, whose behaviorist tendencies lead him to deny, famously, that know how is propositional, still emphasizes the importance of intelligent actions being purposive. Ryle writes:

The clown's trippings and tumblings are the workings of his mind, for they are his jokes; but the visibly similar trippings and tumblings of a clumsy man are not the workings of that man's mind. For he does not trip on purpose. Tripping on purpose is both a bodily and a mental process, but it is not two processes, such as one process of purposing to trip and, as an effect, another process of tripping (1949, 22).

Ryle's position is that instances of know how are intelligent because agents are responsible for regulating and controlling their own performances. That is, they perform certain actions in certain ways and they do this on purpose. Another way of saying this is that their performances are intentional. Of course, Ryle famously denied that being intentional was to be accounted for by some independent, prior, ghostly mental operation of intending that causes the intentional action but his particular way of accounting for intentions has no bearing on the conviction that skill, when instantiated, are intentional in the first place.

Note that in philosophy, it is very hard to find consensus on any issue and the question of whether skills, when instantiated, are intentional actions is no exception. Though most philosophers agree that to perform an action skilfully requires performing it intentionally, Hubert Dreyfus (Dreyfus and Dreyfus 1986; Dreyfus 2002) has claimed that skills, at the pinnacle of expertise, that is, in experiences of flow, do not involve the representation of a goal.[2] Since, Dreyfus grounds this claim in phenomenological evidence of experts who report that in cases of flow, it is as though the action is pulled

[2] For instance, Dreyfus writes: "…in absorbed, skillful coping, I don't need a mental representation of my goal (2002, p.378). Based on this quote alone, Dreyfus may be interpreted as denying only the fact that the intentionality of action requires the representation of a goal but not that expert action is goal-directed or intentional. This would make Dreyfus's position in some ways akin to Ryle's. However, Dreyfus's analogy of intentional action to the forces acting on a soap bubble, where he writes: "Merleau Ponty uses the analogy of the soap bubble. The bubble starts as a deformed film. The bits of soap respond to local forces according to laws that happen to work so as to dispose the entire system to end up as a sphere, but the spherical result does not play a causal role in producing the bubble. The same holds for the final gestalt of body and racket in my example" Dreyfus 2002, p.379) make this charitable interpretation difficult to sustain. It seems more accurate then to ascribe to Dreyfus the position that on his account expert action is not an instance of robust intentional action.

out of them, we should ask, first, whether phenomenology is a good guide to mental metaphysics. Of course, since phenomenology can only tell us about the presence or absence of conscious mental states we have no reason to think that in cases of flow, a goal state is not represented, albeit nonconsciously.

Further, there is positive reason to think that even in cases of flow, intentions guide actions by specifying their ends and the ways in which those ends are to be achieved. Otherwise, there is no explanation for why an agent performs the actions she performs as opposed to other equally afforded actions that are possible in the same environment. After all, there are multiple actions one might undertake in any given situation. Even simple environments have multiple affordances for the same person. Standing at my bathroom sink, I may brush my teeth, or wash my face, or just look in the mirror. And the more complex the environment, presumably, the more affordances it will have. A jungle gym for a child affords climbing, jumping, hanging, sliding, twisting, etc. Determining which of those affordances is acted upon can't be something determined by the environment or by the environment-agent pair in the absence of an agent's goals since, by hypothesis, there are multiple equally afforded and likely actions that the agent can perform.

And this is even more so in well-trained contexts, since most skilled agents will have practiced and routinized a number of actions in the very same circumstances. For instance, a gymnast on beam can perform a full turn, a cat leap, a front handspring or a dismount pass. She has trained a variety of individual tricks and their combinations. The beam together with her physical capacities alone cannot serve to determine which action she will actually perform. She needs to decide, to form a goal, to pursue her intention. She has to move herself towards a specific end.[3]

This point should be even more salient in open sports like soccer or basketball where decisions for what to do are made on the fly, in response to changing circumstances, other players, and opponents. The player must determine, even at seemingly miniscule timescales, whether, e.g., to pass, dribble, or shoot. And part of the agent's skill is constituted by her ability to choose to do the best thing, given the circumstances in which she finds herself. Of course, making this kind of decision need not be like deciding on a career or a life partner—but it is still something that the agent does. Of course, there may be one option that shows up for the agent as obvious—as the only option to pursue—seemingly appearing without deliberation or reflection. All of this is perfectly compatible with the position I am advocating. Given the environment, the agent has many options that she could pursue, and part of her skill is constituted by her ability to "see" exactly which one of those options is best. But this seeing isn't determined for her, rather, it is something she determines by engaging, attending, searching, and playing. In doing this, she sets a goal, intention, or plan according to which her action is to unfold.

Moreover, without a precise plan of action represented by the agent, we have no explanation for why an agent performs her skill in exactly the way that she does.[4] That is, we have no explanation for why the agent pursues an action using the strategy or technique or means that she does, given that, presumably, the more expert the agent the more ways she will have for performing the same task. Of

[3] See Antony (2002) for more on this.
[4] See Fridland (2014) for a development of this point.

course, one might say that the environment directly triggers an entrained, habitual motor routine but then we still have to grapple with the fact that (1) variability is a characteristic of even the most expert performances (Yarrow et al. 2009) and (2) habitual motor routines seem to be the very opposite of skill, especially insofar as flexibility and control are concerned. In fact, compulsive or directly triggered motor behaviours are characteristic of disorder—not expertise. For instance, patients who have suffered lesions in prefrontal areas sometimes exhibit "utilization behaviors" (Lhermitte 1983), where the presence of ordinary objects, such as a water glass or hairbrush, compels the patients to use the objects in the ordinary and familiar way. That is, subjects cannot refrain from performing the familiar, habitual behaviors when in the presence of these objects. It would be a strange account that equated compulsive actions of this kind with the very pinnacle of skilled action.

Taking these considerations into account, a better account of the phenomenon of flow, it seems to me, is that flow experiences are not intention-less but intervention-less. That is, intentions represent planned performances in flow just as they do in ordinary cases. However, these particularly immersive, effortless, and rewarding performances of flow proceed according to routines that go precisely according to plan, and thus need not be constantly corrected, modified and adjusted.[5] That is, agents do not regularly intervene to correct their actions when in flow because the actions do not require the usual kinds of adjustment or modification in order to fulfil the agent's plan. The point, of course, is that a plan or intention is present in flow just as in cases of ordinary skilled and intentional action. Also, of course, present are the requisite mechanisms responsible for monitoring the unfolding action to ensure that it goes according to the plan. This fact can account for the feelings of full absorption and concentration that are characteristic of flow. One important difference between flow and other kinds of skilled performances, then, is the lack of actual need to interfere with the action. And the absence of this need, it seems to me, can account for the effortless and rewarding feelings also characteristic of flow experience. Taken together, we see that even in cases of flow, skills are instances of intentional actions.

To end, it seems to me that a commitment to skilled actions being intentional actions is relatively uncontroversial for good reason. After all, skilled gymnasts do not just find themselves flung into the air, compelled by the presence of a beam or a floor mat to do full twisting double layouts. And surgeons don't usually find their arms magnetically pulled by the various attracting and repelling forces of the operating theatre. Skilled agents perform their skills with every intention of performing them, and sometimes with very particular ideas about the ways in which they intend to perform them. Accordingly, we should accept that the skills are expressions of what agents intend to do. Returning to the interface problem, we should now try to characterize the nature of the intention that is involved in interfacing with the motor representations involved in implementing skilled action.

[5] In this way, Dreyfus may be on the right track when it comes to flow, when he says that "I am in control of my movements in the sense that I can stop doing what I'm doing if I will to do so" (2002, p. 380).

3 Intentions

Even though most philosophers agree that skills, when instantiated, are intentional actions, it will come as no surprise to anyone to learn that philosophers have not always agreed on what it is for something to be an intentional action. For the purposes of specifying the kinds of intentions that are relevant for skilled action, I'll subscribe to the causal theory of intentions that has been forcefully and, to my mind, persuasively forwarded by Michael Bratman (1999). On this theory, intentions are mental states. More specifically, they are plans that specify the actions or outcomes to which the agent that has them is committed. Such plans represent, cause, and explain action. That is, they specify the goal or plan of action, they cause the agent to move towards that goal by performing the planned actions, and they provide reasons for the agent's action. Moreover, intentions are distinctive, *sui generis,* mental states[6]—that is, they are not equivalent or reducible to other mental states such as beliefs or desires and yet they are still subject to strong rationality constraints. That is, intentions are not themselves mental states like beliefs or desires but they are nonetheless personal-level, rational states of the agent. This means that intentions are broadly semantically integrated with the epistemic states of the agent.

Bratman focuses on three main functions of intentions: (i) intentions are the terminators of practical reasoning. This highlights the fact that in a situation where one is deliberating about whether to *a* or *b,* once one has formed the intention to *a,* she no longer has to deliberate about whether to *a* or *b.* Second, intentions prompt further practical reasoning insofar as they create a demand to select the appropriate means for reaching the end that has been settled on with the formation of the intention. So, if one has decided to *a,* one doesn't have to deliberate any longer about whether to *a* or *b,* since that question is settled. But, now, one has to figure out *how to* a. That is, one has to reason about the appropriate means for implementing the plan to *a.* Lastly, Bratman emphasizes that intentions serve a coordinating function—such a function ensures that the agent's actions are coherent over time and also allows cooperation among multiple agents.

Beyond the role of intentions in serving what Elisabeth Pacherie has called, "practical reasoning functions" (Pacherie 2006) outlined above, several philosophers have argued that intentions also play a practical role in governing the online execution of action. In order to serve in this more practical, action-oriented role, it has been suggested that at least some intentions also have the function of (a) initiating, (b) sustaining, (c) guiding and (d) monitoring actions (Mele 1992, 2009; Pacherie 2006, 2008).[7] As Pacherie explains, the differences between the "practical reasoning functions" and the practical "executive functions" of intentions "have led a number of philosophers to develop dual-intention theories of action, or theories that distinguish between two types of intentions involved in action. For instance, Searle (1983) distinguishes between prior intentions and intentions-in-action, Bratman (1987) distinguishes between future-directed and present-directed intentions, Brand (1984) between prospective and immediate intentions, Bach (1978) between intentions and executive

[6] Contra the influential, reductive account of Donald Davidson (2001), which cashes out intention in terms of beliefs and desires or pro-attitudes.

[7] Note that Pacherie thinks that "all seven functions are proper functions of intentions" (2006).

representations, and Mele (1992) between distal and proximal intentions" (Pacherie 2006, 147). To further complicate the vernacular, because my dual-intention model will differ from those already offered, I will refer to action-oriented, executive intentions as "practical intentions" and I call intentions that govern action at somewhat of a distance, "general intentions".

The way the distinction between these two types of intentions has been drawn varies slightly from philosopher to philosopher but the general idea remains the same: some intentions are further removed from action and coordinate plans at larger time scales while other intentions are more practical—guiding and structuring actions as they unfold. The crucial point here is that to implement plans successfully, agents must be receptive to the specific features of their situation and their position in it. But there seems to be no principled or practical reason for all the details required for successful action execution to be incorporated into an agent's general, future-directed, distal plan. So, practical, action-oriented intentions are posited to do the job. Practical intentions coordinate between more general plans and their actual, situated implementation in action.

When it comes to skill, it will be practical intentions that are most relevant for manifesting skilled actions. Though, of course, distal or general intentions will certainly play a role in motivating, organizing, and planning skill, it is practical intentions that are responsible for specifying the precise way that an agent intends that the skill should be instantiated. And since the precise way in which a skill is performed has everything to do with its skilfulness, it is practical intentions that we will need a clear account of. After all, agents at very different levels of skill may share the same general intentions. For example, both a novice and expert skier may have a general intention, while sitting on a ski lift, to ski down the mountain slope, once they reach the top. So, this intention isn't what differentiates the skilled from the unskilled skier. The point that I want to emphasize here is that the difference between the skilled and unskilled agent is not simply that the skilled agent is better able to implement their general plan (which, of course, she is) but that skilled agent is better able to implement that general plan, in part, because she has more appropriate and effective practical intentions about how to do what she plans to do. That is, the more skilled an agent, the more organized and defined her plan for executing a skill will be and, subsequently, the better able she will be at implementing that precise plan in exactly the way that she has cast it.[8] As such, practical intentions are key to an account of skill.

IV. Understanding the differences between practical and general intentions.

In this section, I will spell out what I take to be characteristic of practical intentions and specify how this differs from general intentions. This will further help us to see where intentions need to connect with motor states in order to account for the execution of skilled action. Specifically, I will argue that in order to understand the role of intention in skill, we should draw a distinction between

[8] Notice that nothing about what I've said about practical intentions requires them to be conscious mental states each time they are instantiated in skilled action. I do think that intentions ought to be accessible to skilled agents, in principle, since they are personal-level states that meet strong rationality constraints. It is important to see, however, that their accessibility does not follow from their functions of guiding and sustaining action but, rather, from their integration with the general psychological and epistemic economy of the skilled agent. However, that they are accessible to consciousness, in principle, of course, does not mean that they must be conscious each time they are tokened.

practical and general intentions in terms of their content—more precisely, in terms of which aspect of an action or action plan the content of the intention specifies. General intentions specify ends or outcomes and practical intentions specify means or plans for reaching those ends. The position of an intention in a plan, whether practical or general, depends critically on how the agent conceives of her goals in a specific situation at a particular time.

We should notice that this position differs from Bratman's (1987) dual-intention account, which focuses on differences in timing between future and present-directed intentions in order to distinguish them. And it also differs from Mele (1992, 2009) and Pacherie's (2006) construal of dual-intentions, since I'll deny that general and practical intentions differ in function. In fact, I will maintain that all intentions guide and sustain actions but it is not the job of intentions of any kind to initiate, monitor, or determine all the motor details of an action. Rather, I'll maintain that these latter functions are carried out by control structures that govern the implementation of intentions into actions, but which do not themselves qualify as intentions. I will provide arguments below for why practical intentions are best construed in the way I suggest.

On the view I favour, intentions ought to be conceived of as control structures,[9] which: (1) have the function of *identifying, structuring or specifying* the action(s) to be performed. That is, intentions have the function of *guiding* action insofar as they specify the plan to be executed. (2) As intentions are plans that we are committed to, they have the function of *motivating and sustaining* action. Further, (3) intentions meet rationality constraints. That is, general and practical intentions must be well integrated with the epistemic states of the agent both in terms of their means-ends coherence and their belief-desire consistency. However, I'll maintain that practical and general intentions differ in terms of which part of the action plan they specify. More precisely, general intentions specify the end or outcome of the intention and practical intentions specify the means.

4 Practical and General Intentions: it's Not the Timing

On Bratman's dual intention account, the key feature distinguishing future and present-directed intentions is their timing. Bratman explains that present-directed intentions are concerned with what to do *beginning now* while future-directed intentions concern, as the name suggests, what to do in the future (Bratman 1987, 112). Though the temporal proximity to action is sometimes an important difference between practical and general intentions, it doesn't strike me that timing is key to the conceptual distinction required here.[10] After all, unless we think of general intentions as states that necessarily exist

[9] Note that intentions are a kind of control structure but they are not the only kind of control structure relevant for skilled action. The control structures that are responsible for motor control, monitoring and initiating skilled actions are like the practical intentions that guide skilled action in that all are learned through practice but, because these other functions are not reducible to differences in the content of intentions, it is important to give an account of them that is independent of our account of intentions.

[10] Perhaps, as Pacherie (2006) emphasizes, the timescales at which each intention functions is important for distinguishing distal from proximal intentions—but processing timescales are not identical to temporal sequence. This much should be clear since we might have different intentions that function according to distinct timescale processing principles (e.g., motor processes and propositional states) but all existing simultaneously to govern and structure action.

prior to practical intentions, which then necessarily precede action, we won't get a strict temporal distinction between kinds of intentions.

But it certainly doesn't seem like we have to think of intentions in this temporally sequential way. Instead, we can think of intentions as hierarchically organized, where intentions at higher or more abstract levels of description causally influence, structure, and organize intentions lower down in the intentional hierarchy. This kind of top-down causal influence would allow, and in some cases even require, the simultaneity of intentions at different levels of action specification.[11] So, rather than think of intentions as ceasing to exist once their contents have been handed off to intentions that are temporally more proximal to action, we can think of intentions at more general levels of description (distal or future-directed intentions) exerting top-down influence on more proximal or practical intentions that specify actions at a more fine-grained, contextually grounded level of detail. That is, practical intentions take into account the specific, situated, egocentric, perceptual and bodily affordances of action spaces. In this way, the causal influence of intentions comes through their ability to specify, structure, and shape actions as they are unfolding—not simply prior to their implementation but throughout action instantiation.[12] Accordingly, the distinction between different kinds of intentions will be found in their position in the intentional hierarchy and, respectively, in differences in the detail of content that is specified by them, but not in their timing. But, if this is the way that we ought to characterize the relationship between intentions of different kinds, then timing couldn't be the relevant difference between general and practical intentions.

Applying this insight to skill, we should notice that skilled actions should turn out not simply to be specified by intentions at or prior to initiation but will continue to be sensitive and responsive to intentions at different levels of organization throughout their performance. Moreover, it will be an ongoing communication between intentions and motor representations that will need to be accounted for when solving the interface problem. That is, even after we specify which intentions non-accidentally link up with which motor representations, we will still need to consider how their connection can be continuously maintained and updated throughout performance.

5 Practical Intentions Do Not Initiate Action

On several dual-intention accounts, practical or proximal intentions have been assigned an initiation function. In this section, I will argue that practical intentions do not initiate action even though they do play some role in the initiation of action, namely, that of specifying the conditions under which action should be initiated. In fact, I will claim that it is not the role of intentions at all but rather of initiation control structures to originate action. I will appeal to examples of skill in order to support this claim.

One function that it has been suggested that proximal intentions fulfil is that of initiating action (Mele and William 2009; Pacherie 2006). After all, something has to

[11] This is not to suggest, of course, that temporal sequence has no role in hierarchical organizations but only to highlight the fact that in such an organization, simultaneity or temporal overlap of intentions, is possible.

[12] See Pacherie (2006) and Wu (2013) for detailed discussions of this kind of top-down influence of intentions.

get an action started and since proximal intentions are in the business of specifying and guiding actions as they unfold, it has seemed natural to think that these action-guiding intentions should also be responsible for initiating action. But do intentions really initiate action? It may seem that they could: for instance, think of standing on a 100-m diving board, looking down at the water that appears ever so far away and thinking to yourself, okay, I'm going to jump on 3…2..1! But even in cases that look like intentions explicitly and directly trigger actions, it isn't obvious that initiation mechanisms are short-circuited. After all, isn't the "3, 2, 1" precisely a way of setting the conditions for jumping? And couldn't these conditions be specified by intentions without the action itself being initiated by them? In any case, it should be clear that such direct initiation is hardly the norm.[13] Normally, we just begin acting at more or less the appropriate time—we don't explicitly, intentionally, initiate our action at all. In addition to these logical-space clearing considerations for a contrast between intention and initiation control structures, below, I'll give positive reason to prefer an account that distinguishes between intentional control structures and those that govern action initiation.

To be clear, my claim here is not that intentions play no role in the initiation of action. Rather, I maintain that intentions, general or practical, are not usually themselves responsible for initiating actions. On the view I favour, as action plan specifiers, intentions prescribe the conditions under which an action should begin in more or less detail depending on the circumstances and familiarity of the action. Then, control structures, which are not themselves intentions, initiate the prescribed action once those conditions have been registered as obtaining.[14] On the account I offer, intentions specify the conditions for action initiation but independent initiation mechanisms are responsible for actually initiating action once those conditions have obtained.

To see why the account I favour is a better suited to explain the initiation of action, let's turn to skill. Unsurprisingly, perhaps, the independence of intention and initiation of action seems especially well suited for explaining skill. This is because the initiation of skilled actions are themselves the sort of thing that can be developed and refined independent of the refinement of the content of the relevant practical intention. Think, for instance of high-level, competitive running or swimming where initiating one's actions at the exact right moment has a significant impact on the split-second differences that often mean winning and losing the race. Presumably, all the racers have the intention to start swimming or running, in pretty much the same way, at the exact moment that they hear the gun go off. But actually implementing that intention, that is, initiating the action, can be done more or less effectively, more or less successfully, depending on the skill of the athlete. As such, because the initiation of the action and the intention to initiate the action can come apart in this way, it seems to be that the initiation of the action is not best construed as a function of the intention. Accordingly, I'll hold that it is not the function of intentions to initiate actions but, rather, in the course of specifying a plan, to prescribe the conditions under which the action ought to

[13] As with online motor guidance below (see footnote 19), it may be possible in extraordinary circumstances for intentions to directly cause or initiate action. Think of the case where a victim of stroke is relearning to walk or use a damaged arm. It may be that intentions can bypass ordinary initiation mechanisms in order to directly get action started. The point that I'd like to insist on here is not that intentions can never initiate action but that, in ordinary cases, they do not.

[14] Thanks to Myrto Mylopoulous for very helpful conversations on this point.

be initiated. But it is control structures of another kind, that is, control structures responsible for the initiation of action, that are responsible for actually getting the action started.[15],.[16]

Now, it may be objected that these considerations do not rule out the possibility that intentions act as proximal causes of action initiation but only show that the timing of practical intentions can be better or worse.[17] That is, more expert racers are better able to time their practical intentions, which then initiate their running or swimming actions more expertly. However, in this case, the sufficient condition for initiating action is the mechanism that gives rise to the practical intention and not the intention itself and certainly not its content. In this way, the intention could be the most proximal cause of the initiation of the action but the control of that initiation would still be governed by mechanisms that are not themselves intentions. Either way, we see that it is not the content of the intention that governs initiation of action and it is not the intention itself that determines its own timing and thus the timing of the initiation of action.

6 Intentions Guide Action but So Do Control Structures of the Motor System

Another function that it has seemed natural to attribute to proximal intentions or intentions-in-action is that of guiding action (Brand 1984; Heckhausen and Beckmann 1990; Thalberg 1984, pp. 257–59, Mele (1992), Pacherie (2006). In some ways, the guiding function of intentions is trivial, since intentions—all intentions—represent the action to be performed, all intentions guide action. Accordingly, when practical intentions specify the way the action ought to be undertaken, that is, the structure and organization of the skill, this specification should function as a guide. At the same time, however, it is not obvious that intentions are the only control structures involved in the online guidance of action. In fact, it seems that even though practical intentions guide skill by specifying and structuring the way that a task is to be performed, much of the guidance of the actual implementation of the action will not be done by an intention but by the control structures of the motor system.[18]

When it comes to guiding action, then, it seems that intentions are partially but not fully responsible. To start, it is useful to draw a distinction between guidance in terms of

[15] We should notice that this way of framing things also gives us a straightforward way to counter the challenge from Libet-style experiments, admitting that initiation of action is not something that is done by intentions but still retaining the control and agency that we want over human action.

[16] It also seems that this way of framing things fits more naturally with an evolutionary explanation of action control. Intentions, presumably, could piggyback on primitive mechanisms of action initiation, which should be sensitive to primitive kinds of signals, either environmental or affective. Intentions could then add to the way in which the circumstances for action initiation are specified without thereby duplicating the function of those initiation mechanisms.

[17] Thanks to reviewer 2 for this excellent objection.

[18] On Pacherie's view, low-level motor guidance, which occurs at very short timescales, too fast for conscious control, deliberation and/or access, are performed by motor intentions or M-intentions. Where my account differs from Pacherie's is in distinguishing between intentions and other kinds of control structures that are not themselves intentions but still critically involved in guiding skill. In the coming section on rationality, it will become clear why the motor control structures that I posit are not themselves best conceptualized as instances of intention.

low-level, online correction and modification of the motor and kinematic details of an action and guidance in terms of specifying, structuring and organizing the planned action or outcome. This distinction is the similar to what Pacherie (2006) has called high-level and low-level guidance. I'm happy to follow in this terminology. Of course, intentions guide actions insofar as they specify the plan that is to be executed. What I mean to dispute, however, is that intentions are in the business of correcting, modifying, adjusting, and adapting the kinematic details of an unfolding action. In fact, it would be impossible for intentions, general or practical, to perform this latter online guidance control function since the correction and adjustment of action implementations needs to happen at timescales much quicker than personal-level, deliberate intentions could possibly function. In fact, the basic need for fast, immediate, online correction of action has led to the now widely accepted theory of motor control, which posits forward models at its center (Wolpert et al. 1995). Forward models require that an efferent copy of the expected sensory outcome be sent ahead to a comparator, which can then immediately assess afferent feedback for a match or send a command for appropriate adjustment. That is, even the motor system cannot wait for motor feedback but needs to send an envoy. It should be clear that such online guidance cannot be done by intentions.[19]

Moreover, we see that the more skilled agents are, the more similar their representations of their skills become (Schack and Mechsner 2006). Applying my terminology, we can say that, with expertise, the practical intentions guiding and organizing skill become less variable as skill develops. In fact, we see that not only are novice representations of skill less accurate according to the actual biomechanical requirements of the task, but there is a great deal of variation in novice representations of skill. On the other hand, experts represent their action, the form, structure, order, and organization of the skill, with a great deal of similarity. However, since ability to perform skills still varies between skilled agent's, that is, the ability to control the action details varies even at the highest level of skill, this kind of nuanced action control cannot be reduced to the ways in which the actions are represented. In short, the differences in the actual execution of skilled actions by skilled agents are not reducible to differences in practical intentions.

Of course, having the right intentions is required for having control over an action. The point here is simply that having the right intentions is not sufficient for guiding every part of action as it unfolds. In addition to practical intentions, which are themselves control structures that specify, organize and guide the way in which a skilled action is performed, we also need to posit independent control structures that can handle the fine-grained kinematic details that manifest in motor control. On the account I favour, practical intentions account for the strategic control that is characteristic of skill while control structures of the motor system are responsible for guiding, adjusting, modifying and correcting the low-level motor details of skilful action. Both

[19] Of course, this doesn't entail that intentions or other personal-level states can never directly guide action instantiations. As Briscoe and Schwenkler have argued, during learning or in especially difficult circumstances, it is in fact very likely that intentions take over the guidance of even the fine details of motor control. But these are instances that lack the fluency of normal action and certainly of skilled action. And if we are interested in those cases, which for the purposes of this book we are, then it must be the case that such control is not accounted for by intentions.

kinds of control are acquired through practice and both kinds of control manifest intelligence, though in different ways.

7 Practical Intentions Are Not Monitoring Mechanisms

One further function that has been attributed to practical, intentions-in-action is that of monitoring an unfolding action. This function is closely related to that of guiding insofar as the ability to guide an action is dependent on being able to monitor the progress of the action being executed and judging the extent to which that action conforms with the plan. The idea is that information about how a plan is progressing and to what extent the actions one is executing conform to the plan will be necessary for keeping the plan on track.

Now, I don't disagree that monitoring action's progression relative to the intended plan of action is required for successful execution of intentional action but, again, it doesn't strike me as particularly obvious that monitoring mechanisms should be considered to be functions of intention, general or practical. Presumably, monitoring an action in order to assess its progress and conformity with the intended plan will require a combination of, at least, perception, attention, and understanding. But why should the mechanisms required for attention or understanding be construed as components of intention? After all, we know that visual attention develops with expertise. That is, we see clear, systematic changes in visual attention as a result of sustained, deliberate practice.[20] Now, if monitoring the progress of an action depends on being able to deploy visual attention to evaluate how one is doing, which it seems it must, then unless we say that the development of visual attention in expertise is always reflected in a change to the content of one's practical intention, which isn't at all likely, then it isn't clear why we should think of monitoring as a function of intentions. That is, if the mechanisms involved in monitoring develop independently of the development of one's practical intentions then it isn't clear why we should cash out those mechanisms in terms of intentions rather than, say, as I'll suggest, in terms of attention-related control structures.

After all, keeping track of what one is doing and what one has done is critical for keeping an action plan on track but, as many will find familiar, monitoring, tracking, and evaluating one's progress can be a tricky thing. This is often the case with complicated sequential tasks where one knows the order of steps that one needs to implement but gets lost as to where one is in that order. Think of following a complex recipe, an instruction manual, or a dance routine. Moreover, monitoring and knowing where one is, even in a single movement, is not a straightforward task. In the context of sports practice, for instance, coaches encourage athletes to make changes to specific parts of their movements (think of a tennis swing or yoga pose), but noticing and responding to the precise section of a movement that one is meant to change, is no small feat. That is, it is not only difficult to make the change, it is also difficult to notice, when acting, exactly when and where the change should be made, even when one has been told and understands what needs to be done. The point is that it is often difficult to

[20] See Mann et al. 2007 for a literature review and Fridland (2017) for more on the perceptual and attentional changes that accompany expertise.

notice, to monitor, to attend properly to an action, even if one's intentions regarding the action are skilfully structured. As such, we can see that monitoring one's actions is something that can be improved with practice, something that contributes critically to the successful implementation of action, but it isn't something that is necessarily reflected in the content of the practical intention guiding the action.[21] That is, we can have the same intention and differ in ability to monitor the implementation of that intention. Again, then, we see that the control structures that are involved in monitoring action implementations are not necessarily well accounted for by attributing their function to intentions.

8 Practical Intentions Sustain Action but Not Commitment

Some theorists have claimed that proximal intentions have the function of sustaining action (Mele and William 2009). The thought is that if an agent doesn't actively maintain her intention concerning an action or sequence of actions then her action plan will fail to be performed through to completion. It isn't exactly clear to me why the sustaining function is attributed specifically to practical intentions but I suppose it's not unnatural to think that since it is actual, instantiated actions that need to be sustained, and since practical intentions seem to be in the business of governing actual, occurent actions, then it will be practical intentions that have the role of sustaining these actions, too.[22]

Upon reflection, however, this doesn't seem to be the right way of cutting things. After all, if sustaining an action is, as it seems, a matter of continuing to have the motivation to fulfil one's goals, having the intention should itself be sufficient to sustain action. After all, having an intention is defined as being committed to an action plan. As such, if one continues to have an intention than one should continue to be committed to seeing one's actions through to completion, thus sustaining the action.[23] My general point, then, is that sustaining an action is not something that needs to be accounted for by positing an additional kind of practical or proximal intention that has functions in addition to those of general or distal intentions. That is, intentions are not commitments to plans, the implementation of which need to be sustained by further intentions. Rather, the commitment to the plan is what sustains the actions that one has planned. That is, usually we see an action through to its end in order to fulfil our intentions, simply because we are committed to them (Bratman 1987; Hills 2007). In this way, I take it that I am committed to the standard view of intention here.

[21] See Wu (2013) for more on this.

[22] Notice that Mele (1992) *Springs of Action* considers this to be a function of intentions, generally but in *Effective Intentions* (2009) he says that sustaining actions is something that proximal intentions do.

[23] Notice that this is very similar to Mele's (1992) position concerning intentions. Mele writes: "In addition to any triggering or initiating function that intentions may have in the etiology of intentional action, they also have a sustaining function. This is indicated by what we would quite properly expect to happen if, in the midst of A-ing, an agent were to cease intending to A, namely that he would cease A-ing proceeded so far that it cannot be stopped). Even if the purely representational or nonconative aspect of the former intention were to remain intact (more on this later), we would expect the cessation of intention to bring a halt to the A-ing. This indicates that the sustaining function of intentions, as they are commonly conceived, is at least partly conative or motivational" (130–131). Pacherie (2006) seems also to think that all intentions have all the functions ascribed to P-intentions by Mele (1992) and so sustaining would not be exclusive to P-intentions.

Let's take an example: I have the general intention to do a cartwheel on beam and I have the practical intention to perform this cartwheel very slowly, pausing at the top in a straddled handstand. Now, it seems obvious to me that if I continue to be committed to the general intention, and if the means that I've selected to achieve that general intention remain constant, then there is no reason to think that any additional sustaining of the cartwheel is necessary. That is, as long as my commitment is strong, I'll perform my action all the way through to completion. Importantly, it is not the case that it is only my commitment to the plan of the practical intention that sustains the action. Rather, it is my commitment to the plan, very generally, to intentions of both kinds, that has the function of sustaining action.

Things, however, are never totally straightforward. After all, commitment is a famously fickle thing. However, despite this fact, we should notice that accounting for commitment by positing further intentions will prove to be a non-starter. If we did, we would be off on a regress: requiring a second intention to secure our commitment to the first and a third to fix our commitment to the second and so on and so forth. As such, commitment has to be accounted for in some way other than by positing additional intentions.

Further, we also have the problem that commitment is something that comes in degrees but intentions are not graded states. And accounting for commitment without degrees would have bizarre implications. After all, it would seem strange to have to form a new intention each time our commitment to a plan waivered. For instance, if I wake up in the morning totally 100% committed to eating a healthy lunch and then, by 11 am, feeling a bit peckish, I start picturing a big plate of cheesy pasta—feeling tempted but still *mostly* committed to eating healthy, then, by noon, recovering my sense of resolve and following through on my intention, it seems natural to say that throughout the first part of the day I maintained an intention, though my commitment to the intention waivered a little. It doesn't seem right to say, instead, that I had three independent intentions. But this is what we would have to say if each degree of commitment corresponded to a separate intention. So, if we think of intentions as bivalent, that is, one either has an intention or not, but commitment as graded, then we need to have an independent explanation of the variations in commitment that are not entailed by possession of an intention.

This point becomes all the more striking when we realize that, in reality, the having of an intention only results in fulfilling that intention in about 20–30% of cases (Gollwitzer 1999, 493). Staying committed to the intention requires real work in juggling a host of competing factors. We need to overcome the very real possibility of becoming distracted, we need to prevent other incompatible goals from overshadowing our plan, to persevere through difficult stages of the plan, to reassess and adjust when things don't go according to plan, etc. The real point here is that the abilities involved in maintaining commitment are not the same as those involved in forming or having intentions.

Overcoming temptation, staying focussed, staying disciplined, staying faithful are by themselves the sorts of things that we can do more or less well. That we can get better at. But, it doesn't seem that people have intentions more or less well, though, of course, they can have better or worse intentions. The normativity of the intention does not derive from our ability to sustain our commitment to our plans. It would seem, then, that the commitment one has to an intention requires explanation that is not provided by

an account of possessing intentions alone. One way of dealing with these considerations is to admit that the sustaining of actions is something that commitment to an intention ensures but then explain commitment in terms of other states, processes or mechanisms. That is, we can say that intentions require commitment to a plan but admit that the amount of commitment that one has is not itself something that is accounted for by our theory of intentions.

In this way, we could say that sustaining actions is a function of intentions in general, that is, if one has an intention, the action plan that one is committed to will be seen through to completion. However, we can still maintain that sustaining commitment to the plan requires independent accounting.[24] That is, we can accept that an account of what it takes to maintain commitment to an action plan requires more explanation than an account of how intentions sustain actions, even if sustaining action plans requires ongoing commitment to a plan.

Now, one might think that practical intentions need to have a sustaining function because one is committed to a view where general intentions are temporally prior to practical intentions and cease to exert influence over the action once the appropriate practical intention arises. That is, one may think that general intentions go out of existence once practical intentions begin guiding action. But as I've argued above, this way of thinking about the relationship between different kinds of intentions is overly simplistic. In fact, in order to conceive of the relationship between general and practical intentions properly, we don't need to think of them as occurring sequentially. Rather, we ought to think of intentions as actively exerting a top-down influence throughout the unfolding of an action, especially a skilled action, specifying, structuring, and sustaining the action all the way through to completion.[25]

9 Practical and General Intentions: Means and Ends

I've said a bit about what practical intentions don't do but, as of yet, I haven't provided much in the way of a positive explanation of the role these intentions play in skilled action control and, subsequently, why a distinction between practical and general intentions is worth upholding. First off, it is worth noting that on the account I favour, it is practical intentions that will need to "interface" with motor representations in order to coordinate skilled action. I will provide a defence of this position below by appealing to empirical studies of motor control. For now, let me characterize the way in which I'd like to draw the distinction between practical and general intentions: what really matters for distinguishing general from practical intentions is the role that the content of each kind of intention plays in specifying action. On the description I favour, both general and practical intentions guide and sustain skill but they are responsible for guiding and sustaining different aspects of the skill. General intentions are responsible for providing the end or aim of an action while practical intentions specify the means, way, or technique that the agent will use to achieve that end. Because general and practical

[24] A different approach to accounting for the fluctuations of pro-attitudes that an agent may have toward an action can be seen in Shepherd (2017b) where commitment to an intentional action remains constant while motivation falters (Shepherd gives the poignant example of grading papers). It is certainly possible to cut conative attitudes and thus explanation of graded pro-attitudes in several ways.

[25] See Pacherie (2006) and Wu (2013) for more on how to conceive of top-down influences of intentions.

intentions identify different parts of the skill to be performed, their specifications will require different amounts of detail.

This is not to say that the content of general intentions must always be sparse. To the contrary, an end specified by a general intention can vary in detail. For instance, a general intention can be fairly broad, like "to run a race" or more detailed, like "to run the Boston Marathon in 2020". Still, the practical intention prescribes the means required to successfully instantiate the action. So, in the case of running a marathon, a practical intention will fill in the details for how I can implement the plan—by detailing a schedule of practice, noting the application procedures, finding a way to get to Boston, etc. If all of that has been done, and the Marathon is about to begin, the general intention may remain to run the race and the practical intention would then detail how one would run it, perhaps by specifying the amount of energy one plans to exert or the speed at which one will run at different stages or legs of the race.

The point is that successfully instantiating a skill will require much more than a more or less detailed plan. That is, actually implementing the intention to run a marathon will involve, at the very least, intentionally running, perhaps even in a pretty specific way, in order to avoid aggravating an injury, let's say. It will definitely require running in such a way that one will be able to preserve the right amount of energy to make it to the end of the race. And running isn't just a generic concept that can be used to guide action. The kind of running that needs to be done to successfully run a marathon has to be specified by the marathon runner. After all, intentionally sprinting and intentionally jogging are two different ways of running and one has to decide which one will do and when. But even within jogging, there a million and one variations that one can decide on to guide one's running of the race. The technique can be specified in terms of one's breath or the contraction of one's abdomen or the placement of one's heels on the ground. These specific determinations will then act as guides for how the agent will perform the skill.

Of course, one need not generate a separate conscious intention whose content specifies the exact way that each part of every skill needs to be performed in each instance, one by one, and all the way down. That would be impossible. Some aspects of the skills I have, I perform in such a routinized fashion that even if I used to intend to do them in certain ways, I now contribute very little explicitly to the specification of the manner in which they should proceed. But, we should notice that lacking a conscious or explicit intention of technique does not entail that these aspects of a skilled action are not structured and organized in very particular ways as the result of training. Moreover, these implicitly intended portions of a skill cannot remain unexamined or untouched forever. Even those aspects of a skill that have become fully automatic must be revisited and refined regularly, usually in practice, if one is going to counter arrested development (Sutton et al. 2011; Montero 2016). And this is an important distinction between ordinary action and skill—after all, in ordinary action there is no reason to think that most of the automatic components of our action techniques for, say, walking or eating or typing, will require regular re-evaluation, reorganization, and refinement.

However, we should also note that there are many cases where we as agents do control *the way* or the technique in which a skill plan unfolds, even if we don't control at each moment, each and every aspect of the action. My point here is that many states that are very much under the governance of personal-level intentions are states that specify, with some degree of care, *the way* in which an agent attempts to negotiate her

action space. That is, many personal-level states are in the business of specifying not only ends but means. I call these states practical intentions.

It should be clear that practical intentions must become more specific the closer they get to action. So, whereas, of course, the timing or proximity of intentions to actions matters, it is the role of the content, that is, whether determining means or ends, and the resultant differences in specificity of that content, which actually does most of the work of distinguishing intentions of different kinds. This is a more substantive way of drawing the distinction because it is the content of the intention that has the role of specifying or prescribing what needs to be done and, thus, guiding the action. After all, the content is the plan. Further, we should notice that because in order to implement actions successfully, we will require various states specifying different aspects of a plan with different degrees of detail and at different timescales, there may be several layers of practical intentions governing any one action. That means that there may be several levels of practical intentions that fill in means, and those intentions will vary in how specific their means-specification is.

A natural way to conceive of these layers as being arranged is in hierarchical fashion, each specifying the means of the layer above it. For now, I hope it will suffice to say that I conceive of higher order, more general intentions passing along goal specifications to more practical, particular, grounded, and detailed intentions that are more closely involved in online action control. Specifically, a general intention, which specifies the overall end of an action is inherited by a fairly high-level practical intention, which specifies the means to that end. Next, that practical intention passes down the goal, which is now a single-unit—a means-end unit—to the next and more specific practical intention in the action hierarchy, which then specifies the means for achieving that goal. And so on until the interface with motor representation.

I should add that whether an intention qualifies as general or practical cannot be determined exclusively by looking at the inherent content of the intention but, rather, requires determining the relative position of the intention in the overall plan of the agent. In this way, a practical intention in one situation or time may become a general intention in a different situation or at a different time. As such, whether an intention qualifies as practical or general depends on how the agent conceives of her plan and where the intention falls relative to it. For instance, at time t1, I may have the general intention to make tomato sauce by practically intending to chop, sauté, and simmer tomatoes and onions. At time t2, I may have the general intention to chop tomatoes by practically intending to use a chef's knife and a wooden cutting board and a specific type of dicing technique. As in Aristotle's Nicomachean Ethics, where we see that the ultimate end of all human life is *eudaimonia* or flourishing, we can appreciate that along the way, we may have a multitude of ends or goals that are not identical with this ultimate goal.

10 The Rationality Constraint

At this point, it is worth noting that I am committed to the idea that practical and general intentions not only guide and sustain action but also meet strong rationality requirements. That is, I take it that such intentions are properly cognitive, intelligent, personal-level, agential states that can enter into practical reasoning. In this way, I agree

with B&S in their characterisation of intentions. Moreover, I take it that this characterization of intention is key for a theory of skill because it ensures that skills are the kinds of things that people do, that people are responsible for doing and, as a result of which, various kinds of person-level predicates can be ascribed to them.

Now, Pacherie (2006) has argued that in addition to what she calls distal and proximal intentions there are also motor intentions (M-intentions) that guide action production. These are the same kinds of states that I mean to give an account of in terms of motor control structures. However, there has been some controversy as to whether these kinds of states should really qualify as genuine intentions (Brozzo 2017; Butterfill and Sinigaglia 2014). It strikes me that the main sticking point in this debate is not entirely substantive, however, since both Pacherie and her opponents agree on the characterisation of the motor states in question. That is, they both agree that motor intentions do not meet strong rationality constraints. The question at issue, then, is whether subpersonal, motor, action-guiding informational states should be called "intentions" despite not being personal-level states that participate in practical reason. The main thing to keep in mind, as far as I'm concerned, is that general and practical intentions *can and must* meet robust rationality requirements while motor representations that guide the detailed kinematics of action production *do not.* And, in investigating the interface problem, what we are investigating is how the content of practical intentions hooks up with the content of motor states (whether we call them motor intentions or motor control structures or motor representations).

Before clarifying how I think we ought to understand the rationality constraints on practical and general intentions, it is worth noting that control structures are not differentiated from intentions insofar as rationality is concerned. On the account I favour, practical intentions are themselves a kind of control structure. Practical intentions are responsible for the strategic control characteristic of skill and practical intentions can meet rationality constraints. Other control structures, like those responsible for motor control and attention control, however, cannot meet strong rationality constraints (though they are not excluded from rationality altogether). In short, some control structures like intentions can meet strong rationality constraints while other do not. I prefer to limit the categorization of intention to states that meet strong rationality constraints in order to designate psychological states at the personal-level that straightforwardly participate in practical rationality. It strikes me that this way of cutting things retains that essence of how philosophers usually talk about intentions and, for this reason, avoids some confusion that may otherwise arise. As I said above, however, this decision is largely a semantic issue.

To see exactly what is required for intentions to meet rationality requirements, it is useful to turn to Bratman's theory of intentions. As is well known, Bratman's account of intentions puts a heavy emphasis on rationality. But it is also worth noting that according to Bratman, rationality is a pragmatic notion. He writes, "When we make judgments of agent rationality, then, we apply certain general standards of reasonableness. We require…that the agent's relevant habits, dispositions, and ways at arriving at decisions and actions come up to certain levels of effectiveness in their expected impact on that agent's long-term interest in getting what she wants" (Bratman 1999, 53).

Now, in order to extend rationality into action, Bratman suggests that intentions need to meet two demands. First off, intentions need to be consistent with an agent's beliefs and desires. This requires internal consistency within the plan and also a broader

consistency across an agent's beliefs and desires. Second, Bratman insists that intentions need to be means-end coherent. That is, the way in which specific sub-plans or sub-intentions fill in the action details of more general intentions requires the selection of an appropriate, effective, reasonable means for achieving the end. In short, if intentions are to be successful in coordinating action and reaching the goals of the agent then they need to be consistent with an agent's beliefs and desires and means-end coherent. As Bratman writes, "Both of these demands are rooted in a pragmatic rationale: their satisfaction is normally required for plans to serve well their role in coordinating and controlling conduct" (1999, 31–32). Based on this reasoning, we can see that if a state is both belief-desire consistent and means-end coherent then that state can meet strong rationality requirements and, thus, if that state has the function of specifying an action plan to which an agent is committed, then that state will qualify as either a practical or general intention, depending on the content it specifies.

At this point, it is worth mentioning briefly that in order to observe strong rationality constraints, I am not committed to a view where intentions must be formulated in propositional code. In this way, my approach to the format of intentional content is a bit different from B&S's and more in line with Shepherd (2017a) and Ferretti and Caiani (2018). At root, it seems to me plausible and even likely that we often plan our actions in a distinctly sensorimotor way—generating visual and motor imagery that allows us to weigh different actions and decide on the course of action that we are committed to implementing. That is, we can commit to a plan that is generated in modal or sensory format. This way of perceptual reasoning, that is, agent-level thinking in perception, strikes me as central to action guidance. However, if we insist that thoughts are purely propositional while sensory states are perceptual then these personal-level, motor-modal, intentional states will not qualify as propositional.[26] It is worth noting, however, that even on Bratman's account of pragmatic rationality, there is room for both propositional and nonpropositional states to meet the strong rationality constraints. As Bratmans writes,

[W]eighing conflicting desire-belief considerations…is frequently not a straightforward matter of calculation or deductive inference. In particular, many times the agent's imagination will play complex and important roles in such weighing. The agent will attempt to weigh conflicting reasons by rehearsing in imagination just what would be involved in acting on one or the other of those reasons." (1999, 59).

This kind of reasoning, what Bratman following Dewey calls "dramatic rehearsal", is a kind of imagistic reasoning and planning that is not best construed as logical or syllogistic reasoning where images can function as premises in a deductive argument. Rather, it is a kind of perceptual or imaginative reasoning. But I take it that dramatic rehearsal is a bona fide way of deliberating and deciding on the best course of action. That is, we can commit to plans that are formulated in perceptual code in order to guide our action. Moreover, perceptuo-rational intentions of the kind I have in mind easily interact with what we know, what we believe, what we desire and it influences our decisions and plans in clear, systematic ways. Accordingly, it is possible for practical intentions coded in propositional

[26] Obviously, there are accounts of concepts that take concepts to be or involve sensorimotor elements. If this is the account one prefers, then we can accept that visual or motor states are conceptual. The way I am construing propositions here, however, is according to a distinction familiar in philosophy between thought and perception.

or perceptu-rational format to meet strong rationality requirements because both kinds of states can be widely consistent with beliefs and desires and means-end coherent.

To summarize, I've committed to a view of intentions as distinctive, causal states that specify or prescribe the action plans to which we are committed, thus motivating and sustaining the actions prescribed. However, I've maintained that intentions do not initiate, guide, or monitor action implementations, as others have claimed. Still, I've committed to there being are at least two types of intentions that we need to distinguish between in order to do justice to the control of skilled action: general and practical, where practical intentions will be most relevant to the theory of skilled action. I've said that practical intentions specify the means to the end that is determined by general intentions. I've claimed that practical intentions, in virtue of being closer to action, are more detailed or specific than general intentions and that there can be multiple layers of practical intentions that are arranged hierarchically, where higher level intentions pass along goal specifications to more grounded intentions, which calculate the means to that goal, and then pass that calculation on as a unitized goal to the next intention in the hierarchy. I've also endorsed a view where practical intentions must meet strong rationality constraints and suggested that intentions coded both propositionally or sensorimotorically can do this.

11 Is Practical Intention the Right Level of Intention to Inform Motor Representation?

In this final section, I will review empirical evidence suggesting that motor processing calculates means for an end that is determined, specifically, by the practical intention closest to it in the intentional hierarchy. That is, the goal that the motor system aims towards is neither a low-level sensorimotor goal nor the high level overall purpose of action as specified by a general intention. Rather, the target that the motor system works towards is the instrumental goal as set out at the intermediate level by practical intentions. This seems to be the natural interpretation of the empirical evidence, which shows that motor processing is sensitive to the means that an agent chooses in order to implement action (Mazzoni and Krakauer 2006). Specifically, Day et al. (2016) found that visuomotor learning in a visuomotor rotation task generalizes not around the final target location and not around actual hand location but, rather, around the location where the participant intentionally aims in order to reach the target. That is, the sensorimotor implicit learning of the motor system uses the instrumental goal or what I call the "practical intention" of the agent as its target.

To see why this is the case, let's start with an explanation of the paradigm. The visuomotor adaptation task is often used to examine implicit sensorimotor learning. In this experimental paradigm, a subject's arm is occluded from view and a cursor on a digital screen, which represents her hand location, is used to provide visual feedback. The subject is required to reach towards a target location on a circle and the cursor provides feedback about the success or failure of the reach and sometimes also provides continual feedback about the reaching motion. Experimenters introduce a perturbation by rotating the angle that the cursor moves and, thus, requiring the subject to adjust her hand movement in order to correct for the interference. After repeating the task a number of times, the motor system begins to adapt to the perturbation and correct for it. This adaptation can be observed and

measured clearly with catch trials that capture the after-effect of implicit learning. That is, catch trials that lack perturbation are interspersed in order to establish the extent of adaptation that the motor system has learned. Notably, subjects are usually unaware that any implicit correction or adaptation has occurred (Fourneret and Jeannerod 1998).

We should notice that the visuomotor adaptation task measures very basic motor learning and so it is likely that more sophisticated kinds of motor skill will have more points of contact and integration with explicit, cognitive states. That is, visuomotor adaptation is likely to provide us with evidence of the very minimum level of interaction between intentions and motor representations that is required for task success. Since the motor system is not only a complex system but also a vastly redundant system, it is not the case that all motor learning will be acquired, structured, sequenced or refined in the same way. Still, the visuomotor task is especially useful in that it should provide us with the very minimum threshold of integration needed between explicit and implicit states.

Classically, visuomotor rotation tasks have attempted to avoid or subtract explicit strategies or awareness from their measurements—looking to measure implicit learning exclusively. However, a set of recent studies have sought to capture the interaction of explicit strategy and implicit learning (Mazzoni and Krakauer 2006; Taylor and Ivry 2011; Day et al. 2016; for review see McDougle et al. 2016). This is especially interesting for thinking about the interface problem since such interactions can give us insight into how, at this basic level of motor learning, intentions and lower-level sensorimotor control processes interact.

The first finding that I'd like to discuss is Mazzoni and Krakauer's (2006) study, which introduced an explicit strategy into the visuomotor adaptation task paradigm. In this study, rather than allowing subjects to correct for rotation on their own, participants were given an explicit strategy that they were told to adopt in order to compensate for the interference. That is, Mazzoni and Krakauer told subjects that they should reach at a 45-degree angle in order to correct for the perturbation and thus reach their target. That is, subjects were told that they should aim at 2 o'clock in order to reach the 12 o'clock target. Importantly, using this explicit strategy, subjects were able to eliminate reaching error immediately. That is, the explicit strategy made one-trial learning possible (Mazzoni and Krakauer 2006).

Surprisingly, however, after reaching error was completely and immediately eliminated, implicit adaptation by the sensorimotor system re-emerged making reaching the 12 o'clock target not more accurate but less. That is, the motor system was insensitive to overall task success and continued to adapt according to its own internal processing principles progressively making larger and larger errors. As Mazzoni and Krakauer write,

> Adaptation to visuomotor rotation is independent of explicit strategies, and the motor system overrides an explicit strategy because of the need to resolve conflict between the predictions of a forward model and the visual feedback. Sports coaches should take note that, when it comes to motor learning, the brain has a mind of its own (2006, 3645).

This conclusion, however, is overstated. The problem with this interpretation is not that it posits processing principles that are internal to the motor system and which are distinct from those which govern explicit states. Rather, the problem is that there are actually two explicit target locations involved in this paradigm: one instrumental (to aim at 2 o'clock) and the other ultimate (to reach 12 o'clock). Mazzoni and Krakauer seem only to consider the ultimate target in their interpretation of the relation between implicit learning and explicit strategy. But, importantly, the instrumental target is no less explicit for not being

the ultimate goal of the task. Moreover, as Mazzoni and Krakauer admit, it is the instrumental target that the motor system adopts as its goal. That is, though reaching errors to the ultimate target become larger, reaching errors to the instrumental target become smaller. Accordingly, reaching to the instrumental target becomes more accurate. Mazzoni and Krakauer interpret this as demonstrating an independence between explicit executive states and implicit motor learning but a better way to interpret these results is to specify that the motor system is not sensitive to the ultimate task goal (the general intention) but is sensitive to the explicit, instrumental goal of action (the practical intention). That is, success and failure in reaching the instrumental target are measured in terms of task performance and not simply as a matter of sensorimotor predictions. Of course, this interpretation squares nicely with my position that motor representations interface, specifically, with practical intentions. Moreover, in terms of the rationality constraints on motor structures, we can interpret this finding as suggesting that the motor system is means-ends coherent with practical intentions but not belief-desire consistent with the general psychological states of the agent.

In order to solidify this interpretation, it is also crucial to rule out the possibility that the motor system translates the explicit but instrumental goal into sensorimotor coordinates that are truly independent of the agent's strategy. That is, though we can rule out the motor system's direct sensitivity to the overall, general intention of the agent with the Mazzoni and Krakauer findings discussed above, we still need to show that it is the practical, instrumental intention that the motor system is sensitive to and not a translated, implicit, low-level sensorimotor goal.

The work of Day et al. (2016) is especially suggestive on this front. Day et al. (2016), by examining the generalization of implicit learning around three locations, target (ultimate goal), aim (instrumental goal), and actual hand location (sensorimotor goal), were able to establish that the sensorimotor system is tracking the most frequent location participant's intended to aim at, as it was reported by them. That is, implicit learning generalized around the instrumental goal of the agent, an accessible personal-level goal that subjects had no trouble reporting or verbalizing. But Day et al. were also able to rule out the possibility that the motor system was using a low-level sensorimotor goal as its end, since learning did not generalize around actual hand location.

This study allows us to triangulate on the precise level of intention relevant for guiding the implementation of action by the motor system. That is, it seems that the goal that grounds online motor control is not set by a general intention that represents the ultimate end or overall target of a task nor by a sensorimotor goal specifying location in spatial or sensory rather than semantic terms. Rather, it is the personal-level, instrumental means or way that an agent settles on in order to reach her overall goal that the motor system aims to reach. That is, motor representations are guided towards ends determined by practical intentions.

By considering the nature and function of practical intentions and providing empirical evidence to establish the exact point of contact between intentions and motor representations, I hope that I have made some progress in clarifying how we ought to conceptualize the interface problem. I take it that in getting clear about what it is that intentions do and the points at which they interact with motor processes, we are in a better position to theorize about the kind of communication that we should expect to see at this important juncture between intentions and motor states.

References

Antony, L.M. 2002. How to play the flute: A commentary on Dreyfus's "intelligence without representation". *Phenomenology and the Cognitive Sciences* 1 (4): 395–401.

Bach, K. 1978. A representational theory of action. *Philosophical Studies* 34 (4): 361–379.

Bengson, J., and M. Moffett. 2011. Non-propositional intellectualism. In *Knowing how: Essays on knowledge, mind, and action*, ed. J. Bengson and M. Moffett. Oxford: Oxford University Press.

Brand, M. 1984. *Intending and acting*. Cambridge: MIT Press.

Bratman, M. 1987. *Intentions, plans and practical reasons*. Cambridge: Harvard University Press.

Bratman, M. 1999. *Faces of intention: Selected essays on intention and agency*. New York: Cambridge University Press.

Brozzo, C. 2017. Motor intentions: How intentions and motor representations come together. *Mind & Language* 32 (2): 231–256.

Butterfill, S.A., and C. Sinigaglia. 2014. Intention and motor representation in purposive action. *Philosophy and Phenomenological Research* 88 (1): 119–145.

Davidson, D. (2001). Essays on Actions and Events: Philosophical Essays (Vol. 1). New York: Oxford University Press.

Day, K.A., R.T. Roemmich, J.A. Taylor, and A.J. Bastian. 2016. Visuomotor learning generalizes around the intended movement. *eNeuro 3* (2): ENEURO-0005.

Dreyfus, H. 2002. Intelligence without representation: Merleau Ponty's critique of representation. *Phenomenology and the Cognitive Sciences* 1: 367–383.

Dreyfus, H., and S. Dreyfus. 1986. *Mind over machine*. New York: The Free Press.

Ferretti, G., and Caiani, S. Z. (2018) Solving the interface problem without translation: The same format thesis. *Pacific Philosophical Quarterly*. https://doi.org/10.1111/papq.12243.

Fourneret, P., and M. Jeannerod. 1998. Limited conscious monitoring of motor performance in normal subjects. *Neuropsychologia* 36 (11): 1133–1140.

Fridland, E. 2014. They've lost control: Reflections on skill. *Synthese* 191 (12): 2729–2750.

Fridland, E. 2017. Motor skill and moral virtue. *Royal Institute of Philosophy Supplements* 80: 139–170.

Gollwitzer, P.M. 1999. Implementation intentions: Strong effects of simple plans. *American Psychologist* 54 (7): 493–503.

Heckhausen, H., and Beckmann, J. (1990). Intentional action and action slips. *Psychological review*, 97(1): 36.

Hills, A. 2007. Practical reason, value and action. *Journal of Moral Philosophy* 4 (3): 375–392.

Lhermitte, F. 1983. "Utilisation behaviour" and its relation to lesions of the frontal lobes. *Brain* 106: 237–255.

Mann, D., A.M. Williams, P. Ward, and C.M. Janelle. 2007. Perceptual-cognitive expertise in sport: A meta-analysis. *Journal of Sport & Exercise Psychology* 29: 457–478.

Mazzoni, P., and J.W. Krakauer. 2006. An implicit plan overrides an explicit strategy during visuomotor adaptation. *Journal of Neuroscience* 26 (14): 3642–3645.

McDougle, S.D., R.B. Ivry, and J.A. Taylor. 2016. Taking aim at the cognitive side of learning in sensorimotor adaptation tasks. *Trends in Cognitive Sciences* 20 (7): 535–544.

Mele, A.R. 1992. *Springs of action: Understanding intentional behavior*. Oxford: Oxford University Press.

Mele, A. R., and William, H. (2009). *Effective intentions: The power of conscious will*. Oxford University Press on Demand

Montero, B. 2016. *Thought in action: Expertise and the conscious mind*. Oxford: Oxford University Press.

Noë, A. 2005. Against intellectualism. *Analysis* 65 (288): 278–290.

Pacherie, E. 2006. *Towards a dynamic theory of intentions. Does consciousness cause behavior*, 145–167. Cambridge: MIT Press.

Pacherie, E. 2008. The phenomenology of action: A conceptual framework. *Cognition* 107 (1): 179–217.

Pavese, C. (2018). Know-how, action, and luck. Synthese, 1–23.

Ryle, G. (1949). The concept of mind. London: Hutchinson's University Library.

Schack, T., and F. Mechsner. 2006. Representation of motor skills in human longterm memory. *Neuroscience Letters* 391 (3): 77–81.

Searle, J.R. 1983. *Intentionality: An essay in the philosophy of mind*. Cambridge, MA: Cambridge university press.

Shepherd, J. (2017a). Skilled action and the double life of intention. Philosophy and Phenomenological Research. https://doi.org/10.1111/phpr.12433.

Shepherd, J. (2017b). Halfhearted action and control. *Ergo* 4: 09. https://doi.org/10.3998/ergo.12405314.0004.009.

Stanley, J., and Krakauer, J. (2013). Motor skill depends on knowledge of facts. *Frontiers of Human Neuroscience*. https://doi.org/10.3389/fnhum.2013.0050.

Stanley, J., and Williamson, T. (2001). Knowing how. *The Journal of Philosophy*, 98(8): 411–444.

Sutton, J., D. McIlwain, W. Christensen, and A. Geeves. 2011. Applying intelligence to the reflexes: Embodied skills and habits between Dreyfus and Descartes. *Journal of the British Society for Phenomenology* 42 (1): 78–103.

Taylor, J.A., and R.B. Ivry. 2011. Flexible cognitive strategies during motor learning. *PLoS Computational Biology* 7 (3): e1001096.

Thalberg, I. (1984). Do our intentions cause our intentional actions?. *American Philosophical Quarterly*, 21(3): 249–260.

Wolpert, D.M., Z. Ghahramani, and M.I. Jordan. 1995. An internal model for sensorimotor integration. *Science* 269 (5232): 1880–1882.

Wu, W. 2013. *Mental Action and the Threat of Automaticity*. In Andy Clark, Julian Kiverstein & Tillman Vierkant (eds.), *Decomposing the Will*. Oxford, UK: Oxford University Press. pp. 244-61.

Yarrow, K., P. Brown, and J.W. Krakauer. 2009. Inside the brain of an elite athlete: The neural processes that support high achievement in sports. *Nature Reviews Neuroscience* 10: 585–596.