

# Practical Knowledge\*

*Kieran Setiya*

The phrase “practical knowledge” could mean two different things: the spontaneous “knowledge without observation” that, according to Elizabeth Anscombe and Stuart Hampshire, we have of what we are doing intentionally and, at least sometimes, of what we are going to do; or knowledge how to perform a certain task.<sup>1</sup> This article attempts to restore a commonsense idea: that this is not a mere ambiguity and that these forms of knowledge are intimately connected. It is knowledge how that explains the otherwise troubling possibility of knowledge in intentional action and of the knowledge embodied in our prospective intentions.

The discussion falls into five sections. In the first section, I clarify and briefly defend the idea of knowledge in intentional action. In Section II, I argue that this knowledge is distinctive in being noninferential. The section ends with a puzzle about how such knowledge is possible. This puzzle is refined in Section III through a critique of David Velleman’s well-known theory of knowledge in intention.<sup>2</sup> In Section IV, I argue that our distinctive knowledge of action must be explained through its connection with knowledge how. In Section V, I explore the

\* Some of this material was presented in my fall 2005 seminar on rationality at the University of Pittsburgh; to audiences at Syracuse, Cornell, and Harvard Universities; and at the Wake Forest conference on agency and action in September 2006. Thanks to everyone who commented on those occasions, and in particular to Cian Dorr, Anton Ford, and Doug Lavin. Thanks also to Richard Holton, David Hunter, and Michael Thompson, for written feedback, and to Ned McClennen, who supplied a crucial reference. Finally, I am grateful to Michael Bratman, who replied to the talk at Wake Forest with characteristic generosity and insight; the present version owes a great deal to his objections, encouragement, and advice.

1. See G. E. M. Anscombe, *Intention*, 2nd ed. (Oxford: Blackwell, 1963) and Stuart Hampshire, *Thought and Action* (Notre Dame, IN: University of Notre Dame Press, 1959). I set aside the normative interpretation of “practical knowledge” as knowledge of what one should do.

2. J. David Velleman, *Practical Reflection* (Princeton, NJ: Princeton University Press, 1989).

*Ethics* 118 (April 2008): 388–409

© 2008 by The University of Chicago. All rights reserved. 0014-1704/2008/11803-0002\$10.00

nature of knowing how, arguing, with Gilbert Ryle, that it cannot be reduced to knowing that.<sup>3</sup>

### I. KNOWING WHAT ONE IS DOING

Intentional action is that “to which a certain sense of the question ‘Why?’ is given application; the sense is of course that in which the answer, if positive, gives a reason for acting.”<sup>4</sup> Thus, what we do for reasons, we do intentionally.<sup>5</sup>

On the face of it, intentional action has a second property, that what we do intentionally, we do knowingly. As Hampshire remarks, “[if] a man is doing something without knowing that he is doing it, then it must be true that he is not doing it intentionally.”<sup>6</sup> If I have no idea that I am confusing my audience as I speak, I cannot be confusing them intentionally. And if I am ignorant of the grimace on my face as I ponder an especially difficult question, it too must be unintentional. Once we are clear that this knowledge need not consist in conscious awareness, any more than the knowledge I have of my own name, or of where I am, it is tempting to agree with Hampshire that “doing something . . . intentionally . . . entails knowing what one is doing.”<sup>7</sup>

This conclusion has to be qualified. To begin with, even if it is true that in doing something intentionally one must believe that one is doing it, this belief will not always amount to knowledge. Suppose that I am clenching my fist intentionally, and in doing so I believe that I am clenching my fist. In order to *know* that I am clenching my fist, I must know that I am *able* to clench my fist, in the simple conditional sense: I must know that, if I intend to be clenching it, I will be doing so in fact.<sup>8</sup> Where this knowledge is absent, knowledge of action is undermined. Imagine, for instance, that I have recently been paralyzed and that, with irrational optimism, I believe that I am cured. As it happens, my belief is true: I am able to clench my fist, in the simple conditional sense—though I do not know that I have this ability. If I go on to clench

3. See Gilbert Ryle, “Knowing How and Knowing That,” *Proceedings of the Aristotelian Society* 46 (1946): 1–16, and *The Concept of Mind* (London: Hutchinson, 1949), chap. 2.

4. Anscombe, *Intention*, 9.

5. The converse is less clear, but the dispute about acting intentionally “for no particular reason” will not be relevant here.

6. Hampshire, *Thought and Action*, 95.

7. *Ibid.*, 102. The idea of unconscious knowledge of intentional action may seem odd, but not when we recall the “broadness” of the progressive: even as I sleep, it would be true to say of me, “He is writing a paper on practical knowledge”; I need not be active at the time. This point is emphasized in Kevin Falvey’s “Knowledge in Intention,” *Philosophical Studies* 99 (2000): 21–44, and in Michael Thompson’s “Naïve Action Theory,” forthcoming in his *Life and Action* (Cambridge, MA: Harvard University Press).

8. For related concessions, see Velleman, *Practical Reflection*, 20; and Falvey, “Knowledge in Intention,” 24.

my fist, I may be doing so intentionally, although the belief that I am clenching my fist involved in doing it will not amount to knowledge.

A further complication derives from examples in which an agent acts intentionally in doing  $\phi$  not only without knowledge of what he is doing but without the belief that he is doing it. As Donald Davidson observed, “a man may be making ten carbon copies as he writes, and this may be intentional; yet he may not know that he is; all he knows is that he is trying.”<sup>9</sup> In a case like this, the carbon-copier need not even believe that he is making ten copies, since he doubts that the pressure will go through so many times. The simple belief condition does not apply.

As I have argued elsewhere, the challenge posed by Davidson’s example is limited. Although the carbon-copier does not believe that he is making ten copies, he is doing so by performing other intentional actions of which he *is* aware.<sup>10</sup> For instance, he believes that he is pressing on the paper as hard as he can, and this is the means by which he hopes to make the copies, even if he is not sure that he’ll succeed. We can incorporate this qualification as follows:

If A is doing  $\phi$  intentionally, A believes that he is doing it, or else he is doing  $\phi$  by doing other things, in which he does believe.

A consequence of this principle is that *basic* intentional actions—ones that are not done by doing anything else intentionally—must be accompanied by belief.

There is a final difficulty. What about cases in which the agent’s uncertainty extends further than that of Davidson’s carbon-copier, even to the power to move his own limbs? There seems to be room for a more dramatic failure of belief in intentional action. Consider a variation on the case of recent paralysis in which, at a certain point in my recovery, I am cautiously but not irrationally optimistic: I think that I might be able to clench my fist. Once again, things happen to work out. I clench my fist, and I do so intentionally. Still, given my doubts, I do not believe that I am clenching my fist—perhaps I cannot feel it, or see it, and I am not at all sure of my ability. Nor is it plausible to claim that I am clenching my fist by doing something else, a further intentional action in whose occurrence I do believe. For I deny that I am clenching my fist by an inner act of volition, and I have no thoughts

9. Donald Davidson, “Agency,” reprinted in his *Essays on Actions and Events* (Oxford: Oxford University Press, 1980), 43–61, 50.

10. Kieran Setiya, *Reasons without Rationalism* (Princeton, NJ: Princeton University Press, 2007), 24–25. As Davidson (“Agency,” 50) concedes, when an agent acts intentionally, what he is doing “is known to him under some description.”

about the firing of nerves. There is nothing else that I believe I am doing, as a means to clenching my fist.<sup>11</sup>

Despite appearances, examples like these do not raise a fundamental problem for the view that intentional action is characterized by belief. Instead, they reflect a simplification under which the philosophy of action is typically pursued, akin to the simplification by which epistemology concerns itself with knowledge and justified belief, ignoring the fact that belief comes by degree. A richer description of our psychology would proceed in terms of confidence, or credence, or degree of belief, not just belief and disbelief. And a more sophisticated epistemology would examine the justification of partial belief, not the warrant for belief *simpliciter*. In much the same way, the doctrine that in doing  $\phi$  intentionally one must believe that one is doing  $\phi$  should be qualified not only to accommodate the carbon-copier but also to reflect degrees of belief. If we restrict our attention to basic action, the point is that, in doing  $\phi$  intentionally, one is more confident that one is doing it than one would otherwise be; one has a higher degree of belief. This condition is met even in the case of cautious optimism described above. Although I cannot see or feel my fist, I am more confident that I am clenching it than I was before I began.

Philosophers who doubt the outright belief requirement on intentional action can and should accept this description of the case and the corresponding requirement cast in the psychology of partial belief:

If A is doing  $\phi$  intentionally, A believes that he is doing it or is more confident of this than he would otherwise be, or else he is doing  $\phi$  by doing other things for which that condition holds.<sup>12</sup>

If this principle were false, it would be possible to  $\phi$  intentionally not only without believing that one is doing so but while being quite sure that one is neither doing  $\phi$  nor taking any means to that end; one need not entertain the barest possibility of such things. Beginning with certain

11. This is perhaps too quick. After all, I believe that I am trying to clench my fist, and there are strong arguments for the view that trying to  $\phi$  is a matter of doing something intentionally for the sake of doing  $\phi$ . (See Hugh McCann, "Trying, Paralysis and Volition," reprinted in *The Works of Agency: On Human Action, Will, and Freedom* [Ithaca, NY: Cornell University Press, 1998], 94–109. For a similar view, qualified precisely to omit the attempts of paralyzed agents, see O. R. Jones, "Trying," *Mind* 92 [1983]: 368–85.) If this is right, the case in the text may be one in which I clench my fist by doing something else intentionally, in which I do believe—if only the performance of a mental act. The case would then fall under the reply to Davidson already described.

12. The counterfactual in this principle must be handled carefully. Couldn't there be an action that is normally automatic but which can be done intentionally with a lower chance of success? If one is aware of all this, one will be, on balance, less confident that one is performing that action when one is doing it intentionally—but still more confident than if one were not doing it intentionally and one's automatic system were shut down.

knowledge that one is doing nothing to  $\phi$ , one could do so intentionally without being more confident that one is doing it, or taking means to that end, than one already was. That is hard to accept. Just as one cannot intend to do what one is sure one will not do, so one cannot be doing  $\phi$  intentionally with the unqualified conviction that one is doing no such thing.<sup>13</sup>

The principle of belief-or-confidence that forbids this possibility is sufficient to generate the problems that will occupy us below. For simplicity's sake, I will mostly ignore the complications of credence, adopting the coarser framework of belief and disbelief and the classification of beliefs as knowledge. It would be laborious to translate this into the language of justified confidence throughout. Still, the question that frames the discussion can be stated ecumenically. In light of the first qualification above, the beliefs that figure in doing something intentionally amount to knowledge—or, for partial belief, are epistemically justified—so long as A has knowledge of the relevant ability, in the simple conditional sense. What, if anything, justifies these beliefs? When they amount to knowledge of what we are doing, what kind of knowledge do we have?

## II. KNOWLEDGE IN INTENTION

According to Anscombe's familiar characterization, knowledge of intentional action is "knowledge without observation": when an agent is doing  $\phi$  intentionally, he knows that he is  $\phi$ -ing, and he knows this spontaneously, not on the basis of empirical evidence.<sup>14</sup> She compares the kind of knowledge at issue here with that involved in bodily awareness, the way in which one ordinarily knows, for instance, that one's leg is bent, without having to find out.<sup>15</sup>

As I will argue, Anscombe is right to claim that there is something special about knowledge of what one is doing intentionally; but her explanation is not particularly helpful. For it is not at all obvious how we know the position of our own limbs, or how Anscombe thinks we do. In a later discussion, she observes that we do not infer such facts from "separable" sensations, ones whose content is something other

13. Compare the modest claim that intending to  $\phi$  requires the belief that it is possible to  $\phi$  even if it does not require the belief that what one intends will come to pass (R. J. Wallace, "Normativity, Commitment and Instrumental Reason," *Philosophers' Imprint* 1 [2001]: 1–26, 20). That intending to  $\phi$  conflicts with being certain of the opposite might account for Kavka's well-known "toxin puzzle." (See Gregory Kavka, "The Toxin Puzzle," *Analysis* 42 [1983]: 33–36.) One cannot intend to drink the toxin, however much one would benefit from having that intention, because one knows for sure that one will not act on it when the time to do so comes.

14. Anscombe, *Intention*, 13–15.

15. *Ibid.*, 49–50.

than the proposition inferred.<sup>16</sup> I don't deduce that my leg is bent from feelings of tension or pressure. But, perhaps more strongly, we may not even rely on the vestigial inference that takes us from how things appear to a belief about how they are. For Anscombe, there is "no question of the appearance of the position to me."<sup>17</sup> So, there is an ambiguity here: which characterization is meant to apply to knowledge of intentional action? Matters are further confused by Anscombe's insistence that knowledge of position is speculative or receptive in a way that knowledge of intentional action is not. How does this contrast fit together with her explicit comparison? If we are to make progress here, we will need to clarify Anscombe's remarks.

As I understand it, the crucial claim about the epistemology of knowledge in intentional action is that it is "not the conclusion of an inference."<sup>18</sup> In the typical case of doing  $\phi$  intentionally, I know that I am doing it, and my belief is not inferred from sufficient prior evidence—that is, from evidence that precedes its formation.

Suppose, once again, that I am deliberately clenching my fist. I knew before I began that I am able to clench my fist in the simple conditional sense: if I intend to be clenching my fist, I will be doing so in fact. But this is obviously insufficient to justify the belief that I am clenching my fist. It gives me no reason to believe that I am acting on the ability that it describes or that I intend to do so: it says nothing about what is actually going on. On an inferential model, knowledge of ability is supplemented by knowledge of some other proposition, from which I then infer that I must be clenching my fist. The model can be cast in different forms, depending on the content of the supplement—the further premise to which my inference appeals. I will focus on the defects of two familiar accounts.

According to the first account, knowledge of what I am doing in intentional action, when I have it, is inferred from proprioceptive knowledge of bodily movement. Against this suggestion, Anscombe's objection is sound: "It is an error to push what is known [in intentional action] back and back; first to the bodily movement, then perhaps to the contraction of the muscles, then to the attempt to do the thing, which comes right at the beginning. The only description that I clearly know of what I am doing may be of something that is at a distance from me."<sup>19</sup> In the example she describes, I learn to use a contraption by which an object can be kept level "if I hold a handle and execute a pumping

16. G. E. M. Anscombe, "On Sensations of Position," reprinted in her *Metaphysics and the Philosophy of Mind* (Oxford: Blackwell, 1981), 71–74, 72.

17. *Ibid.*, 73.

18. Hampshire, *Thought and Action*, 70.

19. Anscombe, *Intention*, 53.

movement with my arm and on the downward stroke lower it at the rate at which it would fall.”<sup>20</sup> In keeping the object level intentionally, I know that I am doing so, but I may not know how I am moving my arm, except that I am moving it in such a way as to keep the thing level. Even if I learn how the mechanism works, I am not able to lower my arm deliberately at just the rate at which it would fall. So I have no premise available (through bodily awareness) from which I could responsibly infer that I am keeping the object level. In a case like this, my knowledge of what I am doing cannot be derived from a prior belief about the movement of my body. (A different objection, equally decisive, is that the proposal cannot apply to intentional movements performed under anaesthetic, or to mental actions, like imagination.)

The second account is more tempting, and it has been more influential. According to Paul Grice, knowledge of what I am doing in acting intentionally is inferred from prior knowledge of my will.<sup>21</sup> Here Anscombe’s objection is quite abrupt: “This is a mad account; for the only sense I can give to ‘willing’ is that in which I might stare at something and will it to move.”<sup>22</sup> Maybe so, but it would be nice to have a more persuasive way to make the point. So long as it is possible to have the intention to  $\phi$  without the belief that one is doing  $\phi$ , the proponent of the present view may argue that we can attach an appropriate sense to “acts of will”: they are intentions.<sup>23</sup> Thus, knowledge of what I am doing, when I have it, is inferred from knowledge of what I “will” to do, that is, from what I intend.

The problem with such an account is that it cannot explain why it should be a necessary truth that doing something intentionally is doing it knowingly, in the qualified sense described in Section I.<sup>24</sup> On the inferential model, the connection between intentional action and belief could only be a contingent fact. If, when I am clenching my fist intentionally, the belief that I am doing so had to be inferred from the premise that I intend to be clenching my fist and from empirical knowledge of the conditional that if I intend to do so, I will, it would be possible for the inference not to take place. I might simply fail to put two and two together. Nor is there any guarantee, in such a case, that I am clenching my fist by doing other things, in which I do believe: I need not make

20. *Ibid.*, 54.

21. Paul Grice, “Intention and Uncertainty,” *Proceedings of the British Academy* 57 (1971): 263–79, 278–79.

22. Anscombe, *Intention*, 52.

23. For a detailed theory of intention that might be thought to play this role, see Michael Bratman, *Intention, Plans and Practical Reason* (Cambridge, MA: Harvard University Press, 1987).

24. The strategy of argument here is shared by Setiya, *Reasons without Rationalism*, pt. I, though the focus on epistemology is new.

any inference of the relevant kind. And so I need not have any belief about what I am doing, as I intentionally clench my fist.

We might hope to save the necessity of belief in intentional action, in the face of this omission, by proclaiming that an action is called “intentional” only when it is accompanied by a suitable belief. Although its production is a matter of inference, and is therefore contingent, it is only when the belief is actually produced that one counts as doing  $\phi$  intentionally. But this maneuver is hopeless. It is presumably through the will that reasons are attached to what we do. And on the inferential model, it is possible to will or intend an action, and thus perform it for reasons, without any belief about what one is doing. Since what we do for reasons we do intentionally, it follows that there can be intentional action without belief. And this is the possibility that had to be ruled out.

Anscombe is right, then, to reject the view that knowledge of what one is doing intentionally is derived from prior knowledge of one’s will. Nor would it help to appeal to partial belief. The same objection, about the possible absence of inference, will apply. So long as a gain in confidence about what one is doing is necessarily present in intentional action, its production cannot be a matter of inference from one’s intention. The inferential model cannot account for the fact that beliefs about what one is doing—about the action itself or about the means by which one is doing it—are essential to the operation of the will.<sup>25</sup>

As I have argued elsewhere, the necessary presence of belief in intentional action shows that intention involves belief. This is not a reductive claim: I do not mean that intending to  $\phi$  is simply a matter of believing a certain proposition, since that ignores its distinctive motivational role. But to intend something is, in part, to believe that one is doing it or that one is going to.<sup>26</sup> That is why, as Anscombe pointed

25. The same is true of accounts, like those proposed by George Wilson (“Proximal Practical Foresight,” *Philosophical Studies* 99 [2000]: 3–19, 12–16) and Richard Moran (*Authority and Estrangement* [Princeton, NJ: Princeton University Press, 2001], 124–28), on which knowledge of what one is doing derives from practical judgment, a verdict about the reasons that bear on what to do. The most we can explain in this way is how an agent is “in a position to know” what she is doing, if she takes her action to be determined by the balance of reasons, not why she must believe that she is doing it. Accounts of this kind in any case struggle to accommodate knowledge of intentional action against one’s evaluative beliefs, or when one has sufficient reason for doing more than one thing.

26. For the claim that intention involves beliefs of this kind, see Stuart Hampshire and H. L. A. Hart, “Decision, Intention and Uncertainty,” *Mind* 67 (1958): 1–12; Gilbert Harman, “Practical Reasoning,” reprinted in *The Philosophy of Action*, ed. Alfred Mele (Oxford: Oxford University Press, 1997), 149–77; Velleman, *Practical Reflection*, chap. 4; and Setiya, *Reasons without Rationalism*, pt. 1. Critics of the doctrine include Donald Davidson, “Intending,” reprinted in *Essays on Actions and Events* (Oxford: Oxford University Press, 1980), 83–102, 91–94; Michael Bratman, “Intention and Means-End Reasoning,” *Philo-*



out, the verbal expression of the intention to  $\phi$  is the assertoric utterance of the sentence “I am going to  $\phi$ ” and the expression of “[one’s] intention *in* doing or proposing something” is the assertion that one is doing it.<sup>27</sup> In each case, the expression of intention is also the expression of belief.<sup>28</sup> (These claims must be qualified to deal with partial belief, along the lines suggested at the end of Sec. I. Outright intention involves outright belief; but there is also partial intention, through which one is more confident that one is doing  $\phi$  than one would otherwise be.) It is because intention must be present in intentional action, and because intention involves belief, that beliefs about what one is doing must be present as well. It follows from this conception of intending that knowledge of intentional action could never be inferred from knowledge of one’s intention, since the belief that one is doing  $\phi$  is constituted by the intention one has in doing it; the intention does not precede the relevant belief. This point applies to prospective intentions as well as to intentions in acting. When I intend to  $\phi$ , I often know that I am going to  $\phi$ . If intention involves belief, my knowledge cannot be derived by inference from the condition of my will, which constitutes and does not precede its presence.

There is only one way in which the knowledge that figures in our intentions could be the conclusion of an inference: the intentions themselves would have to be inferred from compelling prior evidence (presumably, in part, psychological evidence) about what we are likely to do. But even if such evidence is always present—as, perhaps, on a form of psychological determinism—this really is a mad account. It implies that we inevitably act under “epistemic compulsion”: on the basis of evidence, not including facts about what we intend, that is sufficient to establish what we are going to do in any case. This conflicts with the experience of decision. Even if it is settled in advance what I am going to do, that is not something I typically know about, in wondering whether to do it. Practical thought is not a search for sufficient evidence from

---

*sophical Review* 90 (1981): 252–65, 254–56; and Alfred Mele, *Springs of Action* (Oxford: Oxford University Press, 1992), chap. 8.

27. Anscombe, *Intention*, 1.

28. Anscombe would resist this way of putting her point: as I understand her view, it is not that intention involves belief but that they are contrasting states of mind. Thus, when “a man is *simply* not doing what he [intends to be doing]”—as in a failure to execute a basic action—“the mistake is not one of judgement but of performance” (Anscombe, *Intention*, 57). On the more natural view, his mistake is one of judgment *and* performance. When I intend to be clenching my fist and I fail to do so, there is a mistake in what I do—but also a mistake in what I believe about myself. For a similar response to Anscombe, see Richard Moran, “Anscombe on ‘Practical Knowledge,’” in *Agency and Action*, ed. John Hyman and Helen Steward (Cambridge: Cambridge University Press, 2004), 43–68, 60–61.

which to predict an action but for considerations on which (for I all know) I may or may not elect to act.<sup>29</sup>

The moral of these arguments is that we have a kind of knowledge of what we are doing intentionally, and of what we are going to do, which is not the conclusion of an inference.<sup>30</sup> It cannot be inferred from bodily movement, since we may not know in relevant, independent detail just how our body is moving—and in the prospective case, it need not be moving at all. It cannot be inferred from facts about our will, since that would make the presence of belief in intentional action contingent, which it is not. Rather, our knowledge of what we are doing, or what we are going to do, is constituted by our will: it is knowledge in intention. And, as experience reveals, our intentions are not inferred from the evidence of action provided by other psychological and circumstantial facts.

The puzzle with which I am concerned below is implicit in these remarks. If they are true, forming an intention is forming a belief about what one is doing, or what one is going to do, but not by inference from sufficient prior evidence. The question is how this transition could be warranted. How can it be reasonable to form such groundless beliefs? As Grice complained in “Intention and Uncertainty,” it is as though having an intention were “a case of licensed wishful thinking.”<sup>31</sup> If anything is not epistemically licensed, if anything counts as a paradigm of defective reasoning, it is believing something you wish were true, or becoming more confident that it is true, without having evidence to show that it is. The exercise of practical reason involved in doing something intentionally or in making a decision for the future can thus begin to seem not only different from but also incompatible with the proper functioning of theoretical reason. In the following section, I press this difficulty through a critical discussion of its treatment in David Velleman’s *Practical Reflection*.

### III. LICENSED WISHFUL THINKING?

Near the beginning of his book, Velleman describes its central project as follows: “I propose to explain [knowledge of action] in a way that accounts for both its timeliness and its spontaneity. I shall explain why, in the normal case, you already know what you’re doing, or at least what you’re trying to do, without ever finding out. . . . My goal is to explain

29. I am not claiming that prediction on the basis of evidence is incompatible with decision (cf. Hampshire and Hart, “Decision, Intention and Uncertainty,” 2), only that it is not required. For related discussion of “epistemic freedom,” see Velleman, *Practical Reflection*, chap. 5.

30. See Hampshire, *Thought and Action*, 70, 128. I reject a final version of the inferential model in Sec. V.

31. Grice, “Intention and Uncertainty,” 268.

these features of an agent's self-knowledge without having to exempt such knowledge from the ordinary requirements of justification."<sup>32</sup>

The initial worry is precisely the one I gestured toward above: at least one ordinary mode of justification is not available here: knowledge in intention cannot justly be inferred from prior evidence. But Velleman denies that this is problematic or that it implies the absence of evidential support. His strategy picks up on a suggestion made by Gilbert Harman in a footnote to "Practical Reasoning": "The problem arises from supposing that the justification of a belief represents a way that someone might reach that belief as a conclusion. But it is also plausible to suppose that one is justified in believing what one believes as long as it coheres with other things that one believes in an explanatorily plausible way."<sup>33</sup>

This is roughly the claim that Velleman makes. Ignoring some irrelevant details of his account, the proposal about prospective intention comes to this: "My answer to [the epistemic puzzle] has two parts. On the one hand, . . . as soon as you [intend] an action, your [belief] is fully justified, by your awareness of [intending] the action and of being inclined to do what you [intend]. . . . On the other hand, the evidence by which you justify your [belief] necessarily includes your having that [intention], and this evidence was lacking until the [intention and belief were] formed. You therefore must have formed the [belief] in the absence of sufficient evidence for it, but . . . you were perfectly within your rights to do so."<sup>34</sup> In other words, it is sufficient to justify forming the belief that one is going to act in a certain way that this belief will be supported by evidence, of which one is aware, once formed.

Suppose, for example, that I decide to make pizza. Before I make the decision, I have no idea what I'm having for dinner and not much evidence one way or another. But having made the decision, I can defend my belief on introspective and empirical grounds: if I know (introspectively) that I intend to make pizza and (empirically) that my intention will be effective, I am justified in believing that this is what I am going to do. According to Velleman, it is sufficient for the *transition* to be justified that it is the forming of a justified belief. Evidence need not precede the belief that it supports.<sup>35</sup>

There are several ways in which this theory will have to be modified or revised. To begin with, it relies on the fact that I not only intend to make pizza but know of my own intention. If the presence of that self-knowledge is contingent, there will be no guarantee that I have the evidence I need. This condition is met by Velleman's claim that inten-

32. Velleman, *Practical Reflection*, 26.

33. Gilbert Harman, "Practical Reasoning," 164 n. 8.

34. Velleman, *Practical Reflection*, 56–57.

35. *Ibid.*, 62–64.

tions are self-referential: they represent themselves as motivating action. So one cannot intend to do something without being aware of one's intention.<sup>36</sup>

More significantly, the theory so far is addressed only to knowledge in prospective intention, knowledge of what I am going to do. What about knowledge of what I am doing intentionally right now? Velleman's answer to this question is not straightforward. He offers two possibilities: first, this knowledge might derive from a suitable prospective intention, one whose content is that I am "just about" to act; second, it might derive from perceptual attention to what I am doing.<sup>37</sup> Neither path is satisfactory, however. The first depicts my knowledge of what I am doing as inferential: it is inferred from prior evidence of what I am about to do. But why must there be such prior knowledge? And what if the inference does not take place? This view cannot accommodate the necessity of knowledge in intentional action. The second path relies on something like the proprioception of bodily movement. It is therefore subject to Anscombe's objection, rehearsed above in Section II.<sup>38</sup>

But none of this matters. Once we abstract from Velleman's interest in the origins of our intentions in acting and focus narrowly on their epistemic side, we can see knowledge of what one is doing as exactly parallel to knowledge of what one is going to do. Like the belief that I am going to make pizza, the belief that I am clenching my fist can be justified introspectively and empirically, so long as I know that I intend to be clenching my fist and that I will be doing it if I so intend. As before, Velleman can simply deny that the evidence precedes the belief

36. *Ibid.*, 88–90, 94–97, 140–41. See also Harman, "Practical Reasoning," sec. 2; John Searle, *Intentionality* (Cambridge: Cambridge University Press, 1983), 83–90; and Setiya, *Reasons without Rationalism*, pt. 1. It is a further question why, and when, this awareness should count as knowledge—but I won't address that here.

37. Velleman, *Practical Reflection*, 51–53.

38. In her revision of Velleman, Hanna Pickard ("Knowledge of Action without Observation," *Proceedings of the Aristotelian Society* 104 [2004]: 203–28) attempts to rehabilitate the appeal to perception, or awareness of action from the inside, at least in explaining how we keep track of what we are doing over time. Part of her argument is that the content of "body awareness" goes beyond such things as movements to include whatever we can do as a basic action (Pickard, "Knowledge of Action," 216–20). Thus, Anscombe's objection does not apply. I am willing to grant the possibility of such awareness and to accept that it plays a role in monitoring the effectiveness of our intentions as we act. But I doubt that knowledge of what I am doing intentionally can originally depend upon it. For, like the inference from prior evidence, the transition from appearance to belief is a contingent fact—and, as I have stressed, the knowledge of what we are doing in basic action is not.

that it supports.<sup>39</sup> In each case, the effect is to embrace a sentiment most famously expressed by William James: “*And where faith in a fact can help create the fact*, that would be an insane logic which should say that faith running ahead of scientific evidence is the ‘lowest kind of immorality’ into which a thinking being can fall.”<sup>40</sup> For Velleman, knowledge in intention falls into the same category as any other kind of self-fulfilling or, better, self-supporting belief: it is a matter of permissible faith.<sup>41</sup>

Unfortunately, I can’t accept this elegant solution to our epistemic puzzle. Like many others, I find the forming of beliefs without prior evidence epistemically suspect, even when the beliefs in question are believed, and even known, to provide sufficient evidence for themselves, once formed. To take an example from another context, consider “the situation you would be in if you believed yourself (perhaps with good reason) to be watched over by a benevolent spirit, who sees to it that

39. A variation on this approach would claim that knowledge in intention is self-verifying in much the same way as Descartes’s *cogito*: if I intend to be doing  $\phi$ , it follows that I am. No doubt this proposal is too simple—it is criticized by Anscombe (*Intention*, 52) and by Keith Donnellan (“Knowing What I Am Doing,” *Journal of Philosophy* 60 [1963]: 401–49, 403)—but there is something to it. The distinction between intending and doing is not a sharp one, and it is tempting to appeal to this in explaining why there is no problem about our knowledge of what we intentionally do. This strategy seems to be invoked by Falvey (“Knowledge in Intention”). He observes that it is normally sufficient to count as doing something that one has started and one intends to go on, regardless of what one is up to right now (Falvey, “Knowledge in Intention,” 25–26). And even quite dramatic mistakes in performance need not undermine my judgment of what I am doing, e.g., I still count as walking home when I am going in the opposite direction after taking a wrong turn (Falvey, “Knowledge in Intention,” 28–29). If the line between intending and doing is blurred, knowledge of action may not be much more problematic than knowledge of intention itself. (A more detailed picture of the close relation between intending and doing is found in Michael Thompson’s “Naïve Action Theory,” which Falvey cites.) Even if we grant this point, however, it cannot amount to a general theory of knowledge in intention. Whatever force there is in the claim that, when I intend to be doing  $\phi$ , I already count as doing it in some incipient way, it does not follow from the fact that I intend to do something tomorrow that I am actually going to do it. The equation of intending and doing—or blurring of the distinction between them—cannot account for knowledge in prospective intention, which does not verify itself. In this respect, at least, Velleman’s theory is more promising than the present approach.

40. William James, “The Will to Believe,” in his *The Will to Believe and Other Essays in Popular Philosophy* (New York: Longman, 1897), 1–31, 25.

41. For this way of putting things, see Rae Langton, “Intention as Faith,” in *Agency and Action*, ed. John Hyman and Helen Steward (Cambridge: Cambridge University Press, 2004), 243–58. I am in broad agreement with her critique of Velleman, which deploys examples like the one that appears in the following paragraph. Two points of dissent: first, I don’t find it helpful to frame the objection by asking whether it is possible to form the relevant beliefs; my focus is narrowly epistemic (cf. Langton, “Intention as Faith,” 255–56); second, I don’t think that the problems for Velleman are problems for the very idea that intention involves belief (cf. Langton, “Intention as Faith,” 252–54).

whenever you form a belief on a certain subject-matter (say, the winners of horse-races), it is true."<sup>42</sup> If I come to believe that Malabar will win, and I know that I have done so, my belief will be supported by sufficient evidence, and it may well count as knowledge. I know my own belief, and I know that the benevolent spirit will make it true. But it still seems to me unreasonable to *form* that belief. And that is just what I am doing when I make a decision, on Velleman's account. For Velleman, forming an intention is epistemically on a par with forming a belief about the winning horse: in each case, I rely on evidence that will not be in place until the belief is formed. If I am justified in forming one belief, by his lights, I am justified in forming both. But then we have to conclude that I am justified in forming neither.<sup>43</sup>

I recognize that this argument will not be convincing to everyone. Velleman, in particular, will insist that it *is* permissible to form a belief about the winning horse without having prior evidence. I hope, however, that the opposing intuition is sufficiently widespread, and sufficiently robust, to motivate the search for an alternative account. The task for that account is to find the epistemically relevant difference between forming an intention, with its corresponding belief, and making a leap of faith that is justified *post hoc*, as in the example given above. In forming an intention, we are somehow exempt from the requirement of prior evidence—as the failure of the inferential model shows—but not because this requirement is never in place. In explaining why this might be, I will diagnose the flaw in Velleman's theory and provide a more abstract, less tendentious objection.

#### IV. KNOWING HOW

Nothing I have said so far implies that the demand for evidence does not apply to knowledge in intention. It may be true that I know what I am doing or what I am going to do only if my belief is supported by evidence, where this is supplied (in the relevant cases) by knowledge of my intention to  $\phi$  and of the fact that I will  $\phi$  if I so intend. It would follow from this, I think, that the transition in which I form my intention and belief is justified only if the belief that is formed will be one for which I have

42. Cian Dorr, "Non-cognitivism and Wishful Thinking," *Nous* 36 (2002): 97–103, 99–100.

43. It is essential to stress here that the benevolent spirit will arrange for Malabar to win only when I *believe* that he will win. Merely picking him as my favorite, or wishing for a victory, is not enough. We should also acknowledge the possibility of "bootstrapping." Suppose I know that, whether doing so is reasonable or not, I am very likely to end up with the belief that Malabar will win. I can then conclude, from the spirit's promise, that this belief is true. In doing so, I form the belief that Malabar will win on the basis of prior evidence. While this is possible, it is not a good model for knowledge in intention: we do not predict our intentions and form them on the basis of those predictions.

appropriate evidence. The problem raised in the previous section is that, even if this condition is necessary, it is not sufficient to justify that transition. Our task is then to say what more should be required.

Alternatively, we may deny that knowledge in intention rests on evidence at all, finding some other role for knowledge of ability. But the problem does not go away. We are not entitled to form just any belief without evidence. What distinguishes forming an intention with its correlative belief from wishful thinking? If not by having evidence, in virtue of what are we epistemically justified in forming such beliefs? I will argue that, whether or not the beliefs that figure in our intentions are supported by evidence once formed, we are justified in forming them without sufficient prior evidence partly through knowledge how.<sup>44</sup>

The idea of a close connection between doing things and knowing how to do them is the basis of Ryle's famous, flawed argument against the "intellectualist legend" that knowledge how can be reduced to knowledge that. The principal later presentation of that argument is in the following passage from *The Concept of Mind*: "The crucial objection to the intellectualist legend is this. The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle."<sup>45</sup> Here the "intellectualist" claims that any exercise of intelligence in action must be preceded by a distinct exercise of intelligence in mental action. It follows, absurdly, that intelligent action can never begin.<sup>46</sup>

The mystery is how we get from this *reductio* to the assertion, on the following page, that "'Intelligent' cannot be defined in terms of 'intellectual' or 'knowing how' in terms of 'knowing *that*'."<sup>47</sup> A connection must somehow be made between the explicit content of the "intellectualist legend" and the nature of knowing how. The most obvious suggestion is this:

1. In doing  $\phi$ , one exercises knowledge how to  $\phi$ .
2. In order to exercise propositional knowledge, one must first consider the relevant proposition.

44. For a similar claim, though without much development, see Anscombe, *Intention*, 88–89.

45. Ryle, *The Concept of Mind*, chap. 2, sec. 3.

46. There is still the possibility of an infinite series of intelligent actions, with no beginning; if each member of the series is more rapid than its successor, the infinity might even occupy a finite time. But this is presumably irrelevant to the exercise of intelligence by creatures like us.

47. *Ibid.*

Suppose, for *reductio*, that knowing how to  $\phi$  can always be identified with propositional knowledge. And consider an instance of action, like clenching my fist. According to premise 1, this action is an exercise of knowing how: I know how to clench my fist, and I rely on this knowledge in doing so. By hypothesis, there is a proposition, that  $p$ , such that knowing how to clench my fist is knowing that  $p$ . It follows by premise 2 that, in order to exercise my knowledge that  $p$  in clenching my fist, I must first consider the proposition that  $p$ . And now we are embarked upon Ryle's regress. For, in considering the proposition that  $p$ , I must exercise knowledge how to consider that proposition. This in turn consists in knowing that  $q$ . And so I must first consider the proposition that  $q$ . Even if the propositions are not distinct, every act of consideration must be preceded by another—and so it is impossible ever to begin.<sup>48</sup>

This is not, I think, a compelling argument. The first premise is tempting only when doing  $\phi$  is doing  $\phi$  intentionally. I can do all sorts of things that I do not do intentionally without knowing how to do them: growing older and forgetting where I left my keys, for instance. The problem is that, on the corresponding interpretation, the second premise can be denied. In order to exercise propositional knowledge in action—as in being moved by it—one need not intentionally consider the relevant proposition. So the iteration or regress is blocked at the first step.<sup>49</sup> Nor does it help to insist that premise 2 is concerned with the intentional exercise of propositional knowledge, which does require a prior intentional act of considering a proposition. For, on the corresponding interpretation, premise 1 can be denied: in doing  $\phi$ , one need not intentionally exercise one's knowledge how to  $\phi$ . And so the argument still goes wrong.<sup>50</sup>

What interests me here, however, is not the failure of Ryle's argument but the insight embodied in its first premise when this is properly understood. A more plausible formulation would restrict the demand for knowing how to intentional action:

If A is doing  $\phi$  intentionally, then A knows how to  $\phi$ .

48. Compare the version of Ryle's argument discussed by Jason Stanley and Timothy Williamson ("Knowing How," *Journal of Philosophy* 98 [2001]: 411–44, 412–16). They rely on a nontemporal analogue of premise 2, which removes the threat of regress. Instead, they assume that distinct actions correspond to distinct propositions and that "it is [not] necessary to contemplate an infinite number of distinct propositions [in order to engage in action]" (Stanley and Williamson, "Knowing How," 414). In my view, the reading in the text comes closer to Ryle's intentions.

49. For this objection, see Stanley and Williamson, "Knowing How," 414–15.

50. That the exercise of knowing how involved in intentional action need not itself be intentional can be seen from the fact that, in, say, clenching my fist intentionally, I may not realize that I am exercising knowledge how. Thus the belief condition, from Sec. I, may not be met.



Unfortunately, the truth is more complex. There are cases of intentional action that are not accompanied by knowledge how. For instance: I am trying to defuse a bomb, staring with confusion at an array of colored wires. Which one to cut? In desperation, not having a clue what the wires do, whether they will trigger the bomb or not, I disconnect the red wire—and the timer stops. Even though I did not know how to defuse the bomb, and managed to do so through dumb luck, I count as having defused the bomb intentionally. That is certainly what I meant to do, despite my uncertainty. As with Davidson's example in Section I, however, the challenge here is limited. When I do something intentionally that I do not know how to do, I must at least know how to take some relevant means. In the present case, I know how to cut the red wire, and I think it might defuse the bomb, even though I can't be sure. Knowledge how belongs at the core of any intentional action. We can state the connection as follows:

(K) If A is doing  $\phi$  intentionally, A knows how to  $\phi$ , or else he is doing it by doing other things that he knows how to do.

It follows that basic intentional action, not performed by doing anything else intentionally, must involve the exercise of knowledge how.

The need to account for K is a constraint on theories of agency akin to the belief constraint of Section I. Just as it is impossible to perform an intentional action in ignorance—both of what one is doing and of the actions by which one is doing it—so it is impossible to do something intentionally without knowing how to do it or how to take the relevant means. It is this connection that generates the systematic objection to Velleman's theory promised above. In Velleman's account, an agent who is doing  $\phi$  intentionally—and has knowledge in intention that he is doing so—must satisfy only two requirements specific to doing  $\phi$ : he must know that he intends to be doing it, and he must have knowledge of ability, in the simple conditional sense. If this picture is to be consistent with the truth of K, knowledge how must be contained implicitly within these elements. In particular, since they are directed at doing  $\phi$  itself—which is the object of intention and ability—not at the means by which it is done, they must contain the agent's knowledge how to  $\phi$ . The problem is to say where this could be.

It is clear from the very beginning that we cannot identify knowledge of one's intention with knowledge how. Knowing how to do something is a persisting state, not something present only when it is exercised, and the foolhardy may intend to be doing things they do not know how to do. More promising, perhaps, is the view that knowledge how is present in Velleman's conditions as the capacity for intentional action. To say that one knows how to  $\phi$ , on this conception, is just to say that one is capable of doing it intentionally, in the modest sense of

“capacity” in which, whatever one does, one has the capacity to do. The mystery for Velleman is why this capacity, even granting its presence, deserves to be thought of as a form of knowledge. Why think of it as *knowing* how?<sup>51</sup> On his account, the capacity to  $\phi$  intentionally is epistemically inert: all the work is done by knowledge of intention and ability. This capacity may figure in the object of that knowledge, but that does not make it “epistemic”; the proposition that  $p$  is not an epistemic proposition just in virtue of being known. It is a condition of adequacy on any account of K to explain the sense in which knowing how is *knowledge*.

What we are looking for, then, is a state that figures among the requirements of intentional action, that can precede the occurrence of action as a standing condition, and that deserves to be conceived in epistemic terms. (The precise relationship of knowing how to other forms of knowledge is at this point undetermined: it may consist in knowing a proposition or it may count as knowledge in some other way.) In the story that Velleman tells, there is just one candidate to play this role: knowledge of ability must constitute knowledge how. The problem is that, even for Velleman, knowledge of ability is not required for intentional action, only for knowledge in intention. So this proposal could not possibly explain the truth of K. Think back to the example from Section I, in which I hold the irrational true belief that I am cured of my paralysis. I manage to clench my fist intentionally without having knowledge of ability. But I still know how to clench my fist, and I exercise this knowledge in doing it. Knowledge of ability is not required for knowing how.

It follows that something is missing from Velleman’s theory of intentional action. It ignores the need for knowledge how, as an epistemic

51. As Edward Craig (*Knowledge and the State of Nature* [Oxford: Oxford University Press, 1990], 150–53) points out, use of the same verb for knowing how and knowing that is not peculiar to English: it is not an accident that we characterize knowing how in epistemic terms. (I should stress, however, that I do mean to place much weight on linguistic evidence. That the concept of knowledge how is epistemic is something we can tell directly in grasping that concept, just as we can tell that the concepts of warrant and justified inference belong with knowledge, as epistemic, even though we express them with different words.) Craig’s own account of knowing how is not persuasive. He argues that the capacity to  $\phi$  intentionally counts as knowledge because those who possess it tend to be good instructors, in much the way that those who possess propositional knowledge tend to be good informants (Craig, *Knowledge and the State of Nature*, 156). The analogy is defective in that, for basic action in particular, knowing how has little pedagogical significance.

state distinct from knowledge of ability.<sup>52</sup> This objection does not rest on intuitions about the example in Section III: it is based on the need to explain why knowing how is a condition of intentional action and how it amounts to knowledge. Velleman's theory cannot meet these demands. In the final section, I sketch an account of knowledge how on which it is defined by its role in dynamic epistemology and on which we can thereby explain the necessary truth of K. It is a consequence of this account that knowing how to  $\phi$  is not propositional knowledge.

## V. KNOWING HOW AND KNOWING THAT

We have been occupied with two puzzles about practical knowledge: how we can be justified in forming the beliefs that figure in our intentions and how to explain the role of knowing how in intentional action. In each case, a problem was posed for Velleman's epistemology. The problems were independent, but it is natural to connect them. On this hypothesis, it is knowledge how that solves the epistemic puzzle from Section III: it exempts us from, or qualifies, the demand for prior evidence which would otherwise prohibit us from deciding to act. We are entitled to form the beliefs that figure in doing  $\phi$  intentionally only when and because we know how to perform the relevant acts.

This point about the epistemic need for knowledge how applies to prospective intention as well as to our intentions in acting. If I don't know how to  $\phi$ , I am not entitled to form the intention and belief that I am going to do it: knowledge how must supplement knowledge of ability in the epistemic justification of decision. We must not be misled, here, by the fact that we can learn new skills. Can't I decide to dance the tango at my wedding, one might ask, even if I don't yet know how? The answer is that this decision would not be justified. Rather, I must decide to learn how to dance the tango and to exercise this knowledge at my wedding, once it has been acquired. These are things I do know how to do. I can then infer from the knowledge that figures in my intention that I am going to dance the tango at my wedding. But this

52. It is worth noting that the distinction goes both ways. The case in the previous paragraph illustrates the possibility of knowing how to  $\phi$  without having knowledge of ability. One could also have knowledge of ability, in the simple conditional sense, without knowing how. Suppose that, although I don't know how to twitch my nose, I am watched over by a guardian angel who would provide me with that knowledge instantaneously, were I to form the appropriate intention. I might then know that I am able to twitch my nose, in that I would be doing so if I intended to be, before I know how to do it. If we shift to knowledge of ability in the prospective sense—knowing that, if I intend to  $\phi$ , I will do so in the future—the examples are more mundane. I may know that I will execute my intention by learning how to do something that I do not yet know how to do. As I argue below, however, forming an intention to  $\phi$  when one does not yet know how to  $\phi$  is epistemically problematic.

is a prediction, not the content of a decision in its own right. I cannot decide to dance the tango at my wedding without an unjustified leap of faith.

These observations are the key not only to the puzzle from Section III but also to the nature of knowledge how. They show the need to recognize its epistemic role. Knowing how to  $\phi$  is the state or condition that, with knowledge of ability, provides the epistemic warrant for decision. Together, they justify the transition in which one forms the intention and belief that one is doing  $\phi$  or that one is going to do it. More carefully, this transition is justified if and only if one's decision is an exercise of knowledge how to  $\phi$  and one has knowledge of ability, in the simple conditional sense. Knowledge how thus plays a role in dynamic epistemology, in our entitlement to form and revise beliefs. This happens continuously in the performance of intentional action. When I exercise my knowledge how to clench my fist and I have knowledge of ability, I know that I am doing so. As knowledge of the present, this must be renewed as time goes by. In effect, I have to form the belief that I am clenching my fist *now* at each new moment. So my knowing how to do so is constantly implicated in knowledge of what I am doing.

None of this conflicts with the natural thought that, in some sense, knowing how to  $\phi$  is having the capacity to  $\phi$  intentionally by forming and acting on the relevant intentions. The problem is not that this is false but that it is not illuminating. The sort of capacity involved in knowing how to  $\phi$  is hard to specify. In particular, it is not a matter of ability, in the simple conditional sense. I still know how to move my arm when it is tied to my chair, even though intending to move it would achieve nothing. Whatever capacity I have, in knowing how to move my arm, it is not one that I am able to exercise at the time. Nor does the capacity account, by itself, explain why knowing how should count as knowledge. (That was part of the case against Velleman in Sec. IV.) In order to comprehend the epistemic character of knowing how, we need the claims that were made in the previous paragraph. Knowledge how is required to justify us in forming intentions, with their corresponding beliefs, and thus to warrant the persisting beliefs about what we are doing that figure in acting intentionally. (The need to understand why knowing how is a kind of knowledge also provides support for the argument of Secs. I and II. It is the existence of knowledge in intention that makes epistemic work for knowing how in the performance of intentional action.)

Finally, unlike the capacity theory on its own, these claims explain why knowing how is not reducible to knowing that. On the present conception, a decision is epistemically justified by knowledge of ability, together with knowledge how. If knowledge how to  $\phi$  is knowledge of the truth of some proposition, this is an instance of justification by prior

beliefs. Thus, on a propositional interpretation of knowing how, our model of knowledge in intention is inevitably inferential: the belief that I am doing  $\phi$ , or that I am going to do it, that figures in my intention is justified by inference from things I previously knew. The problem is that, as we saw in Section II, the inferential model is false.

It is true that, in arguing against this model, I did not explicitly examine an inference that appeals to the propositional knowledge that constitutes knowing how. But this won't help. There is a special difficulty in relying on knowledge how as the basis for an inference to intention, which turns on the following principle:

If I am justified in forming the belief that  $p$  on the basis of knowing that  $q$ , then the fact that  $q$  is evidence that  $p$ .

According to the propositional account, I am justified in inferring that I am doing  $\phi$  now or that I am going to  $\phi$  tomorrow from the propositional content of knowing how to  $\phi$  together with knowledge of ability. It follows by the principle above that knowledge how to  $\phi$  and knowledge of ability together provide some evidence that I am doing it now and that I am going to do it tomorrow. And this is surely not the case. Knowledge how to do something and knowledge that I am able do not provide me with standing evidence that I am doing it or that I will ever do it again.<sup>53</sup> The epistemic job that is done by knowing how could not be done by knowledge of a proposition. Knowing how is not reducible to knowing that.<sup>54</sup>

The broader picture that emerges from this account is one in which we are sometimes, but not always, permitted to form beliefs without sufficient prior evidence. I am justified in coming to believe that I am clenching my fist, when that is what I intend to be doing, so long as I do so as an exercise of knowledge how and I know that I am able to clench my fist, in the simple conditional sense. This is so despite the fact that I lack sufficient evidence for my belief, at least until it has been acquired. It does not follow, however, that I am entitled to make a leap of faith whenever I know that, in doing so, I will form a true belief. This raises the difficult questions of when we are entitled to form beliefs without sufficient prior evidence and why. It is arguable that we are

53. This appears in the fact that these conditions permit us to form the beliefs in question only by forming an intention. They do not warrant those beliefs per se. And so they make an epistemic discrimination—between intention and mere belief—that propositional knowledge cannot make.

54. This argument applies even to sophisticated propositional views, like that of Stanley and Williamson ("Knowing How"). They identify knowing how to  $\phi$ , roughly, with propositional knowledge of the means by which one could  $\phi$ , under a "practical" mode of presentation. A further argument against this view is that it cannot account for the necessity of K: knowledge of means is not required for basic intentional action.

permitted to do so in cases that do not involve intentional action in any ordinary sense: we need to make room for noninferential knowledge of mental states like beliefs and desires. What is much less clear is how such cases relate to the ones addressed above. Can they be subsumed by the exercise of knowledge how—as, perhaps, through Ryle’s suggestion that we know how to reason or by the association of self-knowledge with mental action?<sup>55</sup> Or do they provide a further and distinct exception to the demand for prior evidence? I cannot pursue these questions here; I mention them mainly as a plea for further reflection. The epistemology of knowing how must find its place in a theory of dynamic justification that tells us when, and why, these exceptions should be made.

55. Ryle, “Knowing How and Knowing That,” 6. Something like the latter association is found in Richard Moran’s *Authority and Estrangement*.