

### III—THE EPISTEMIC ROLE OF INTENTIONS

JOHANNES ROESSLER

According to David Velleman, it is part of the ‘commonsense psychology’ of intentional agency that an agent can know what she will do without relying on evidence, in virtue of intending to do it. My question is how this claim is to be interpreted and defended. I argue that the answer turns on the commonsense conception of calculative practical reasoning, and the link between such reasoning and warranted claims to knowledge. I also consider the implications of this argument for Velleman’s project of vindicating the commonsense view by showing it to be consistent with an ‘evidentialist’ epistemology.

#### I

There is a sharp divide, in the recent literature, between two attitudes to first-person knowledge of intentional actions. Some philosophers take it to be a mark of such knowledge that it requires no evidential support. Others regard it as a condition of adequacy on a philosophical theory of such knowledge that it should identify the evidence on which it rests. A suggestive way to make sense of these attitudes has been proposed by David Velleman. The idea that intending to do something typically provides for ‘groundless’ knowledge of what one will intentionally do or is doing, Velleman suggests, is a central element of the ‘commonsense psychology’ of agency. We are indebted to Anscombe (1957) and Hampshire (1965) for articulating that idea. But we are philosophers as well as commonsense psychologists, and one question we need to address as philosophers is whether our putative knowledge in this area ‘is worthy of the name’ (Velleman 1989, p. 105). If an affirmative answer to that question is to be defended, we need to establish that, and how, our beliefs in this area meet ‘the usual requirements of evidential support’ (1989, p. 25) A central aim of Velleman’s own well-known theory of practical self-knowledge is to show that the sense in which, from a naive standpoint, we take such knowledge to be

groundless is consistent with the sense in which, as philosophers, we can make it intelligible in terms of its (highly distinctive) ‘evidential support’.

In this paper I shall not be much concerned with the details of Velleman’s account of how the ‘manifest’ and the ‘philosophical image’ of practical self-knowledge fit together. My question is whether we should accept his starting point, the idea that there *is* a ‘manifest image’ of such knowledge, that is, a pre-theoretical understanding reflective agents have of the connection between intentions and ‘groundless’ knowledge. You might say that the idea is fanciful, given that ordinary members of the public usually take no interest in how they know what they are doing. Or you might insist that reflective agents are in fact familiar with some evidential basis, for example, a combination of introspection and observation. Some of Velleman’s remarks suggest that he takes his claims about the manifest image to be borne out by reflection on the phenomenology of agency. In acting intentionally, he maintains, you ‘feel as if’ you know what you are doing without observation or evidence (Velleman 1989, p. 20). I will set this line of argument to one side. I want to focus on a different, and I think more promising, kind of consideration, which has to do with our reaction to a certain kind of challenge. Suppose you declare you will attend the staff meeting tomorrow, and suppose someone were to challenge you to corroborate your claim. ‘How can you *tell*?’, they ask, or ‘What makes you think you will attend the meeting?’ Velleman makes two points about this kind of example. One is that you, or any reflective agent, would find these questions not just surprising but somehow improper. You would insist that you are in a position to know without *having* to ‘tell’ (1989, p. 21). The second point is that what explains and rationalizes this response is that you subscribe to a certain explanatory schema: ‘the reason’ an agent knows her actions without evidence, according to common sense, is ‘that the action was her idea to begin with: it’s what she had in mind to do’ (1989, p. 24).

Let me put the upshot of this formally:

- (CP) According to ‘commonsense psychology’, intending to  $\phi$  normally puts one in a position to know one will  $\phi$  without relying on prior evidence that one will  $\phi$ .

By formulating CP in this way I am simplifying things in two respects. First, I shall only be concerned with the case of knowledge of

what one will do, not of what one is currently doing. Velleman, I think rightly, takes an explanation of the former to have an important bearing on that of the latter, but the relationship between the two cases raises complex issues, for example, to do with the various roles perception plays in controlling and monitoring bodily intentional actions, issues I will not consider here. Second, in focusing on CP I'm going to ignore a stronger thesis that features in Velleman's discussion. Common sense, he maintains, credits agents with knowledge that is *spontaneous*, that is, 'generated from within' (Velleman 1989, p. 47), where this in turn is glossed in terms of the thought that one acquires such knowledge by 'inventing' an 'idea' of what one will do. This material is not easy to interpret, and it is harder still to find any support for it. I suspect it is best seen as a description of how practical self-knowledge 'appears to us' in terms that foreshadow the philosophical *explanation* of that appearance Velleman will eventually construct, specifically his analysis of intention as a kind of belief, and of practical reasoning as a matter (roughly) of 'jumping to conclusions'. CP offers a less loaded description of the putative appearance.

At this point it is useful to distinguish three kinds of questions: First, is CP defensible? Second, how is the kind of explanation of self-knowledge CP attributes to commonsense psychology actually supposed to work? Third, is the view of self-knowledge CP ascribes to commonsense defensible? The second question, I shall argue, holds the key to the other two questions. Velleman takes a notably austere view of the explanatory resources of commonsense psychology. Our pre-theoretical understanding, he thinks, is only concerned with how self-knowledge is *acquired*, not with whether what we have is actually *knowledge* (is 'worthy of the name'). Commonsense psychology is restricted to the 'cognitive psychology' of self-knowledge, and has nothing to say about the 'epistemological' question of how claims to knowledge in this area are to be defended. This account of the division of labour between common sense and philosophy shapes Velleman's approach to the third question. Whether the commonsense view is defensible, he thinks, depends on what the correct epistemological theory has to say about the evidential basis of practical self-knowledge. On the theory Velleman develops, the letter (though perhaps not the spirit) of the naive view's denial that our knowledge is based on prior evidence turns out to be correct: practical self-knowledge, it transpires, 'rests on' evidence from

which it could not have been ‘derived’.

Now, while Velleman’s austere answer to the second question dovetails nicely with his conception of the project of assessing the commonsense view, it causes trouble with the first question. The case for CP, recall, crucially depends on the idea that our pre-theoretical understanding of self-knowledge explains *and rationalizes* our rejection of the request for evidence. But it is hard to see how the commonsense view can make it reasonable for us think we are able to know without evidence if it has nothing to say on how a claim to knowledge is to be corroborated. In the next section, I elaborate the challenge, and offer some reasons to doubt it can be met so long as we accept Velleman’s austere analysis of the commonsense view. In §III, I sketch an alternative, less austere answer to the second question, which can be extracted from Stuart Hampshire’s discussion of self-knowledge in *Freedom of the Individual* (1965). The basic idea is that the naive view of practical self-knowledge is inseparable from the naive view of practical reasoning: we think of practical reasoning as something that can warrant claims to knowledge. On this analysis, I will suggest, CP has much plausibility. The problem, from Velleman’s perspective, is that it now looks as if commonsense psychology and an evidentialist epistemology offer rival accounts of what makes our knowledge worthy of the name. I conclude by considering how the conflict might be adjudicated.

## II

Consider first Velleman’s charge that Anscombe and Hampshire are guilty of a non sequitur. They held, plausibly enough, that practical self-knowledge is not derived from prior evidence, but mistakenly concluded that therefore such knowledge ‘doesn’t require evidential grounds’ (Velleman 1989, p. 25) The problem with this inference, in Velleman’s view, is not just that it is rendered invalid by the possibility that knowledge may ‘rest on’ evidence from which it is not ‘derived’. Rather he takes the inference to involve a kind of category mistake. It moves from a claim of ‘cognitive psychology’, about the way self-knowledge is acquired, to a conclusion that ‘belongs to epistemology’, ‘since it’s about how self-knowledge is to be justified’ (1989, p. 25). In Velleman’s striking phrase, Anscombe and Hampshire ‘try to save the commonsense psychology of self-knowledge by

elevating it to the status of epistemology' (1989, p. 25).

One might object that, in general, information on how someone's knowledge that *p* was acquired (e.g. by observation or inference) has an utterly immediate bearing on how the claim that she knows that *p* is to be defended or justified. But of course Velleman may argue that the case of practical self-knowledge is an exception to that rule. Common sense, he might insist, realizes that it is somehow because of your intention to go to the staff meeting that you know you will go, but it has no idea of how to support or defend the claim that you know you will go.

The basic problem with this picture is that on the face of it, it is incompatible with the consideration Velleman uses to introduce and motivate the commonsense view in the first place. For the question 'How can you tell you will go to the meeting?', or 'What makes you think you will go?', is not a question raised from the detached perspective of cognitive psychology. It is, at least potentially, what Austin called a 'pointed' question, bringing into play the possibility that you don't actually know what you claim to know and requesting a reason to think that you do know (Austin 1961, p. 78). Correlatively, if you dismiss the question by insisting that you know without any *need* for evidence, you are insisting that your not relying on evidence does not impugn the credentials of your claim to knowledge. In Velleman's terms, your response 'belongs to epistemology'. We can reinforce this point by considering what you would say if your pre-theoretical understanding really were limited to the concerns of 'cognitive psychology'. In that case, you should only be expected to say that as a matter of fact you consulted no evidence in arriving at your knowledge that you will come to the meeting. The obvious rejoinder would be that you had better start consulting evidence now, or else you should withdraw your claim to knowledge. At this point, you would have nothing sensible to say in response; for by stipulation you have no *reason* to think that you are in a position to know without relying on evidence.

So there seems to be a mismatch between the *resources* attributed to the commonsense view by Velleman and the *work* it is expected to do, a mismatch that threatens to undermine Velleman's argument for CP. As far as I can see, the most promising way to get around this problem, consistent with Velleman's project, would be to reinterpret the cognitive psychology point along functionalist lines. Commonsense psychology, it might be said, has a naive theory of in-

tentions, involving numerous generalizations concerning the typical causes and typical effects intentions have, conditional on the agent's other mental states. One such generalization is that an intention to  $\phi$  tends to go with, or give rise to, 'groundless' knowledge that one will  $\phi$ . Our grasp of that generalization could make it rational for us to attribute knowledge to intentional agents, and to repudiate the request for evidence. On the other hand, this does not imply that we have any understanding of the *reason why* intentions provide for knowledge. To get to the bottom of why the generalization holds, one would need to discover the realizer state that occupies the characteristic functional role of intentions. Velleman's epistemological theory—his view of intentions as beliefs that are justified partly through being self-fulfilling—might be understood as an attempt to identify the psychological mechanisms realizing the commonsense theory of intentional, autonomous agency.<sup>1</sup>

But the functionalist interpretation is not a satisfactory solution. How are we supposed to have acquired the critical generalization linking intentions and knowledge? Presumably, as with the 'theory' theory in the philosophy of mind, the idea would be that our naive theory is supported by evidence, in so far as it enables us to explain and predict observable phenomena. But then once again the commonsense theory would have to be informed by concerns that Velleman thinks properly belong to epistemology: evidence supporting the generalization would have to be evidence that the attitude usually associated with intentions is indeed *knowledge*.

I would like to speculate a bit on how Velleman ended up with his strange picture of the commonsense view as folk cognitive science. I want to highlight two factors, to do with Velleman's reductive aspirations in both epistemology and the philosophy of action. First, there is his entrenched commitment to an 'evidentialist' epistemology, which may disable him from contemplating the possibility that anyone might find the credentials of a knowledge claim intelligible other than by displaying the evidential support enjoyed by the relevant belief. If common sense has no insight into the evidential basis of practical self-knowledge, this would mean it can have no inkling

---

<sup>1</sup> This reading accords with Velleman's general description of his project as that of constructing a 'hypothetical theory' designed to 'identify ... a kind of personal constitution that would exemplify [put differently: realize] our conception of autonomy' (1989, p. 7). It would also be consistent with Velleman's later claim that his hypothetical theory finds support from empirical research in social psychology (2007, pp. xvii–xix).

of how such knowledge is ‘worthy of the name’. The second factor lies in Velleman’s reductive view of practical reasoning as a ‘mode of theoretical reasoning’ (1989, p. 90), involving the formation of expectations about one’s actions by ‘jumping to conclusions’ on the basis of (insufficient) evidence regarding one’s motives, expectations that are said to constitute intentions in virtue of their distinctive functional role. This is not how we ordinarily think about practical reasoning and intentions, nor does Velleman claim it is. What he assumes, though, is that we can articulate the commonsense psychology of practical self-knowledge without paying heed to the commonsense view of practical *reasoning*. (One important difference between the two, from his perspective, is that only the former is susceptible to a (partial) philosophical vindication.) I think this assumption makes it unsurprising that Velleman ends up with a picture on which common sense has no idea as to *why* the generalization linking intentions and knowledge holds. For it is arguably practical reasoning, as ordinarily conceived, that makes it intelligible to us why the generalization holds: our conception of practical reasoning, as it were, presents us with the categorical basis of intentions’ tendency to generate knowledge of one’s intentional actions. I think an extremely helpful, and rather neglected, way to spell out this thought can be found in Hampshire’s analysis of the epistemic role of intentions. In the next section I will suggest that Hampshire’s account yields a plausible defence of CP. (I will leave open the interesting question to what extent Hampshire’s account differs from Anscombe’s.) In the last section, I consider where this leaves Velleman’s distinction between the naive and the philosophical perspective on self-knowledge.

### III

Hampshire’s discussion (1965) revolves around a distinctive kind of statement and the basis that licenses the statement. The statement is ‘I shall come to the meeting’. Its putative basis consists of the propositions ‘It would be a great mistake for me not to attend the meeting’ and ‘I will be able to go or not, as I choose’. I first discuss Hampshire’s central thesis before looking at the naive epistemology of self-knowledge that can be derived from it.

Let me introduce the thesis with the help of an analogy. It is some-

times said that it is a ‘structural element’ of practical reasoning that a judgement that one has most reason to  $\phi$  licenses the formation of an intention to  $\phi$ , with no need for an extra mediating ‘step’ consisting of an expression of the desire to do what one takes oneself to have most reason to do. The idea here is that to demand such an extra step would be to misrepresent something that belongs to the structure of practical reasoning as an additional input or premiss for such reasoning.<sup>2</sup> In that way, the demand would be reminiscent of Lewis Carroll’s Tortoise’s view that, crudely put, we have no reason to accept the conclusion of a sound deductive argument unless an explicit statement of the relevant rule of inference is added to the premisses. Hampshire’s thesis can be formulated in similar terms. It is a ‘structural element’ of practical reasoning, on his view, that the judgements that I have most reason to attend the meeting and that I’m able to go, as I choose, can license the statement ‘I shall come to the meeting’. A salient feature of that statement is that it has a ‘double aspect’ (Hampshire 1965, p. 72). In making it one aims to express both one’s intention to attend the meeting and knowledge that one will attend the meeting. So the ‘structural element’ point means that practical reasoning can warrant not only the formation of an intention but at the same time what Hampshire calls a ‘statement of fact about the future’. The latter does not depend on any intermediate ‘steps’, such as the judgement that one intends to go to the meeting and an estimation of the likelihood of one’s implementing that intention. In that sense, the transition between practical reasoning and the factual statement can be immediate.

Hampshire’s account of the import and the basis of the statement ‘I shall come to the meeting’ distinguishes his position from each to two opposed views on practical self-knowledge. On the one hand, he disagrees with those who think that first-person expressions of intentions should be understood, not as reports or factual statements, but as something akin to orders. On Hampshire’s view, first- and third-person versions of ‘x will attend the meeting’ have this in common: it is a necessary condition of the statement’s being correct that it should express knowledge of the future.<sup>3</sup> (This does not entail

---

<sup>2</sup> For an illuminating exposition and defence of this suggestion, see Stroud (2003). The bearing of Lewis Carroll’s paradox on practical reasoning and action explanation is also a central theme of Stroud (2011, ch. 4).

<sup>3</sup> This interpretation is encouraged by Hampshire’s characterization of statements about one’s future actions as ‘claims to knowledge’, and as being ‘normally required to serve’ the



that if the statement turns out to be incorrect one can properly be blamed for having made it. And it is of course consistent with allowing that its incorrectness may reflect what Anscombe calls a mistake in the performance.<sup>4</sup>) On the other hand, Hampshire rejects the view that first- and third-person versions of the statement express the same *kind* of knowledge, namely, knowledge based on inference from the agent's current intentions. On the rejected view, the two aspects of the statement 'I will come to the meeting' correspond to two different kinds of considerations on which the statement must be *based*: qua expression of intention, it should be based on practical reasoning; qua 'statement of fact', it should be based on one's assessment of the relevant evidence. For Hampshire, a single basis—practical reasoning—can be sufficient.

The disagreement between Hampshire and his inferentialist opponents is unlikely to be resolved on phenomenological grounds. Inferentialists might concede that deliberation issuing in 'double aspect' statements does not usually fall into distinct parts or phases, devoted, respectively, to the practical question of what to do and the factual question of what one will do.<sup>5</sup> They may acknowledge that when answering questions about my future intentional actions 'my concern [is] not with what I should say but with what I should do' (Hampshire 1965, p. 69). Still, they will question whether practical deliberation, even practical deliberation leading to the formation of an intention, can *warrant* a factual claim. After all, it is surely always an open question whether the intention supported by the deliberation will in fact be implemented. In response, defenders of Hampshire's 'structural thesis' should argue that it is precisely *because* the deliberation is practical that it warrants a factual conclusion. For practical deliberation, on their view, should not be mistaken for theoretical deliberation aimed to discover what someone, or oneself, should do. Rather it is deliberation that has a practi-

---

purpose of informing 'others of the probable future course of events' (Hampshire 1965, pp. 58–9). In a similar vein, Broome (2009) argues that 'the inaudible marker for intention does not cancel the assertoric effect of the sentence', in contradistinction to the 'inaudible markers for commands and questions'.

<sup>4</sup> See Anscombe (1957, p. 82). Anscombe's wording ('the mistake') encourages the assumption that there will be only one kind of mistake here. As Pears (1975) points out, that assumption is not mandatory.

<sup>5</sup> See Paul (2009) for a version of an inferentialist account on which the 'inference' grounding our beliefs about what we will do is held to be a matter of sub-personal information-processing rather than conscious reasoning.

cal goal: in the most basic case, it aims to find a means by which one will be able to attain some specific end, and the point of the deliberation itself is to attain that very end.<sup>6</sup> (Thus the deliberation shares its objective with the action in which it issues.) Now any uncertainty as to whether one *will* act as intended will have an immediate practical significance: it would reflect a lack of confidence, either about one's practical ability to perform the intended action or about one's commitment to pursuing the end in question. Either way, the uncertainty would compromise the success of one's practical deliberation. Correlatively, in so far as one has good reason to take oneself to have deliberated successfully, one has good reason to regard the question of what one will do as settled.

We can spell out the issue between the two positions by looking at examples inferentialists might invoke in support of their view. Suppose you form the intention to move a log blocking your driveway, but are not convinced you will succeed.<sup>7</sup> Evidently you do not think your practical reasoning warrants the claim that you will move the log, and surely you may be right about this. If you know you are liable to bouts of forgetfulness, you may not even take yourself to be entitled to claim that you will *try* to move the log. Cases such as these may be thought to reveal a general gap between intentions and warranted 'statements of fact about the future'. They may encourage not just a conclusion in the philosophy of action, that intending to  $\phi$  does not entail believing one will  $\phi$ , but also an epistemological conclusion, that there is a gap between practical reasoning and beliefs about what one will do.<sup>8</sup>

It is this latter conclusion that defenders of Hampshire's thesis want to resist. They need not deny that examples of 'agnostic intentions' are perfectly coherent and can play a significant role in work on the nature of intention. What they deny is that such intentions are prototypical. Note, first, that the subject does not, in these cases, know or even believe that she will be able to perform the intend-

---

<sup>6</sup> See Vogler (2002) for a rich exposition and defence of the 'calculative' conception of practical reasoning.

<sup>7</sup> See Bratman (1987) and Holton (2009, ch. 1), for discussion. Compare also Anscombe's remark that 'the less normal it would be to take the achievement of the objective to be a matter of course, the more the objective gets expressed *only* by "in order to"'. E.g. "I am going to London in order to make my uncle change his will"; not "I am making my uncle change his will" (Anscombe 1957, p. 40).

<sup>8</sup> Michael Bratman appears to draw both conclusions: see Bratman (1987, pp. 37–8).

ed action as she chooses. So the cases present no formal counter-examples to Hampshire's thesis. Still, it might be said that what the examples bring out is that this kind of knowledge is not essential for the purpose of practical deliberation. But note, second, that in cases of agnostic intentions, it is not clear, at least to the agent, that her practical deliberation has been successful. For all you know, you may not have succeeded in your aim to find a means that will enable you to remove the log. This is why it is natural to expect that you will *go on deliberating* about alternative arrangements. (Bratman himself suggests that you should, in addition, form the conditional intention to call in the tree company in case you fail to move the log.) Note, furthermore, that even your agnostic intention involves many things you do know you will be able to do, such as leaving the house, walking to the drive, and so on. In fact, it is not easy to envisage practical reasoning that does not draw on knowledge about one's practical abilities at any point.<sup>9</sup> In so far as it does, one might insist, Hampshire's structural thesis holds. If you have reason to take yourself to have deliberated successfully, you have reason to regard the factual question as settled. To refuse to accept the relevant 'statement of fact', in such cases, would amount to a confusion about the structure of practical reasoning akin to the Tortoise's confusion.

Turning from practical reasoning to epistemology, how would Hampshire's structural thesis—supposing it can be defended—bear on the explanation of self-knowledge? Well, if practical considerations can warrant a 'double aspect' statement, and the norm of correctness for such statements (under one of their aspects) is to express knowledge,<sup>10</sup> then your practical reasoning can make it intelligible how you are a position to know what you will do. One might be tempted to put this by saying that you gain knowledge of what you will do *by* expressing your intentions.<sup>11</sup> The problem is

---

<sup>9</sup> See Setiya (2008) for discussion of the role of 'knowledge how' in making practical self-knowledge possible, though in the framework of a reductive analysis of intention as a kind of belief.

<sup>10</sup> See Williams (2000). As noted earlier (n.4), Hampshire's account is in line with this view.

<sup>11</sup> See Taylor (1985). One might put this idea by saying that we use a certain procedure to gain knowledge of our future intentional actions: we answer the question of what to do, and then, in a separate step, endorse or accept an expression of the resulting intention. This would be one way to understand Falvey's suggestion that we have a 'general warrant' to 'employ' or 'present' expressions of intentions as descriptions of what we are doing, or will do (Falvey 2000, pp. 23, 37).

that this makes it sound as if you do after all use a means—a ‘method of discovery’—for gaining self-knowledge. This is not Hampshire’s view. As he often stresses, the question to which the ‘double aspect’ statement is the answer is a practical question of what to do: in the normal case, ‘my concern [is] not with what I should say but with what I should do’ (Hampshire 1965, p. 69). So I will stick with the formulation that successful practical reasoning ‘puts the agent in a position’ to know what she will do. Importantly, on Hampshire’s analysis, in some sense the connection between practical reasoning and warranted claims to knowledge is part of our ordinary, pre-theoretical view of practical reasoning and intentional action. Even a non-philosopher who declares she will come to the meeting may find her possession of the knowledge expressed in her statement unmysterious in the light of the practical reasoning on which the statement is based. This should not be taken to mean that she has a good answer to the question ‘*How do you know* you’ll come to the meeting?’ For that question is naturally heard as a question about the means by which she obtained her knowledge; and as indicated, in the normal case she will not have employed any means. Still, it is clear to her, given her awareness of her grounds for the statement, that she did not make her statement irresponsibly. It is not, for example, as if the statement were merely a matter of speculation or wishful thinking. On the one hand, then, she is aware of having good grounds for her statement; on the other hand, clearly her grounds are not adequate evidence for the truth of the statement. (Of course evidence may play a role in establishing whether she will be able to go to the meeting, but her judgement that she will go is based on practical reasons.) These points provide the agent with a robust response to the question featuring in Velleman’s argument for CP, ‘How can you tell you will go to the meeting?’: they give her a *reason* to insist that she is in a position to know without ‘having to tell’. Thus Hampshire’s analysis may succeed where Velleman’s failed, in providing an account of how our pre-theoretical understanding of practical self-knowledge *rationalizes* the rejection of the request for evidence. In this way, the analysis may help to sustain the case for CP, hence to secure the starting point of Velleman’s theory. At the same time, it is incompatible with a basic commitment of that theory, that the only way to determine whether our knowledge is ‘worthy of the name’ is to show that our beliefs in this area meet ‘the usual requirements of evidential support’. To adjudicate this

conflict, we need to return to our third question: whether the view CP attributes to common sense is defensible.

#### IV

A question that is conspicuous by its absence from Hampshire's discussion is why he *believes* that he will attend the meeting. The focus of his reflections throughout is on the warrant he has for his *statement* 'I shall attend the meeting'. The warrant bears on the epistemology of self-knowledge, given that the statement aims to express knowledge. But Hampshire does not offer what might be called a belief-based explanation of his knowledge; he does not undertake to validate the claim that he knows he will attend the meeting by showing that his belief satisfies certain general conditions for knowledge, such as enjoying 'evidential support'. It is here that a defender of Velleman's evidentialism may want to step in. She will deny that in the absence of an explanation of how practical reasoning yields relevant evidence there can be any *intelligible* link between practical reasoning and one's possession of propositional knowledge. Of course, she might grant that Hampshire's 'structural thesis' correctly articulates the commonsense view. But she would argue that in that case the commonsense view must be rejected.

Note that Hampshire need not deny that knowing entails believing, or that believing satisfies certain general conditions, such as epistemic justification. The issue is not whether knowledge involves justified belief, but whether an *explanation* of how one is in a position to know that *p* must be underpinned by an explanation of the evidential basis of one's belief that *p*. Note, further, that familiar 'internalist' intuitions do not speak to that issue. It is sometimes said that forming a belief in the absence of adequate evidence is 'epistemically suspect' (Setiya 2008). Yet it can hardly be suspect to believe that *p* when one knows that *p* and has good reason to think one knows that *p*. 'Internalist' intuitions may tell against certain kinds of belief-based explanations of our possession of knowledge, but they cannot tell us whether a belief-based explanation is mandatory.

The issue between Hampshire and the evidentialist turns, I think, on the question of what should count as a reason to accept a particular validating explanation of knowledge. How can we determine whether a putative account of how we are in a position to know

certain facts—an account that would give us a reason to think that the attitudes it makes intelligible are indeed a case of knowledge—passes muster? According to the evidentialist, the claim that *X* is a source of knowledge is defensible only if, in the light of our prior understanding of the general conditions for knowledge, we can satisfy ourselves that *X* generates knowledge by recognizing that it provides for the satisfaction of just those conditions. Hampshire's account does not conform to that ideal. He makes no attempt to justify the claim that practical reasoning yields knowledge by appeal to the way it provides for the satisfaction of the general conditions on knowledge, and it is hard to see how such an attempt might succeed. From Velleman's evidentialist perspective, this amounts to shirking the philosophical question of whether our knowledge in this area is 'worthy of the name'.

The dialectical situation is complicated by the fact that, as we saw, on Hampshire's analysis the commonsense view itself has something to say about what makes our knowledge 'worthy of the name': sound practical deliberation provides us with grounds to affirm that we know what we will do, without having to 'tell'. This suggests the beginnings of a response to the evidentialist challenge. If practical deliberation gives us grounds to rebuff the 'first-order' request for evidence, we might appeal to those very grounds to reject the request for the kind of philosophical vindication envisaged by the evidentialist. It is after all a substantive philosophical claim that any good explanation of how we are in a position to know certain kinds of facts depends on a prior, independent explanation of how our beliefs in this area meet the requirement of 'evidential support'. If the claim turns out to be incompatible with our entitlement to make 'double aspect' statements on the basis of sound practical deliberation, we might argue, this would give us a good reason to reject it.

As it stands, this response will cut no ice with an uncompromising evidentialist, who may be happy to jettison the naive view of practical reasoning along with the folk epistemology of practical self-knowledge. But one might try to develop the response, by showing that a commitment to the naive view is implicit in something the evidentialist will find hard not to acknowledge. One question that might be raised here is whether it is possible consistently to acknowledge that people sometimes act intentionally, while denying that reason can be practical, or at least that there is such a thing as

practical reasoning as ordinarily conceived. To act intentionally, the claim would be, is after all to act for practical reasons (or at least in a way intelligible in terms of the agents taking something to be a practical reason), and our conception of practical reasons is inseparable from our conception of the norms governing the episodes of reasoning in which such reasons are assessed and recognized. Obviously in pursuing this line of argument one would need to address large, controversial issues about the nature of action explanation. I mention it here merely as an illustrative example of an alternative approach to our third question. The project would not be to offer a philosophical defence of the commonsense view of practical self-knowledge, but to question whether it can be coherently challenged.<sup>12</sup> It might be argued, in this way, that a commitment to the naive explanatory schema is integral to our view of ourselves and others as intentional agents, and to our grasp of what it is to act intentionally.<sup>13</sup>

*Department of Philosophy*  
*University of Warwick*  
*Coventry CV4 7AL*  
*UK*  
*j.roessler@warwick.ac.uk*

#### REFERENCES

- Anscombe, G. E. M. 1957: *Intention*. Oxford: Blackwell.
- Austin, J. L. 1961: 'Other Minds'. In his *Philosophical Papers*, ed. J. Urmson and G. Warnock, pp. 76–116. Oxford: Clarendon Press.
- Bratman, Michael 1987: *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Broome, John 2009: 'The Unity of Reasoning'. In Simon Robertson et al. (eds.), *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford: Oxford University Press.
- Falvey, Kevin 2000: 'Knowledge in Intention'. *Philosophical Studies*, 99, pp. 21–44.

---

<sup>12</sup> The difference between these projects is a major theme in the work of Barry Stroud. See, for example, Stroud (2011, ch. 5).

<sup>13</sup> Many thanks to Sarah Broadie, Quassim Cassam, Jennifer Hornsby, Matt Soteriou and David Owens for helpful comments.

- Hampshire, Stuart 1965: *Freedom of the Individual*. London: Chatto and Windus.
- Holton, Richard 2009: *Willing, Wanting, Waiting*. Oxford: Clarendon Press.
- Paul, Sarah 2009: 'How We Know What We're Doing'. *Philosophers' Imprint*, 9, pp. 1–24.
- Pears, David 1975: 'Wanting and Intending'. In David F. Pears, *Questions in the Philosophy of Mind*, pp. 39–55. London: Duckworth.
- Setiya, Kieran 2008: 'Practical Knowledge'. *Ethics*, 118, pp. 388–409.
- Stroud, Barry 2011: *Engagement and Metaphysical Dissatisfaction*. Oxford: Oxford University Press.
- Stroud, Sarah 2003: 'Weakness of Will and Practical Judgement'. In Sarah Stroud and Christine Tappolet (eds.), *Weakness of Will and Practical Irrationality*. Oxford: Oxford University Press.
- Taylor, Charles 1985: 'Hegel's Philosophy of Mind'. In his *Human Agency and Language: Philosophical Papers 1*, pp. 77–96. Cambridge: Cambridge University Press.
- Velleman, J. David 1989: *Practical Reflection*. Princeton, NJ: Princeton University Press.
- 2007: Introduction to the David Hume Series Edition of *Practical Reflection*. Stanford, CA: CSCI Publications.
- Vogler, Candace 2002: *Reasonably Vicious*. Cambridge, MA: Harvard University Press.
- Williamson, Timothy 2000: *Knowledge and its Limits*. Oxford: Oxford University Press.