

## Intention, Self-Governance, and Self-Knowledge

Matthew Soteriou

Talk about what one intends to do is normally reserved for one's attitude toward *future* conduct – what one *will* do. According to Bratman, although this doesn't show that there are not present-directed intentions (intentions directed at what one is *now* doing), it does suggest that in order to understand what intentions are, we should be focusing on the future-directed cases. Over a number of years Bratman has been arguing that future-directed intentions are typically elements in larger plans for future conduct that help us to coordinate and project our agency in an organised way over time. According to Bratman, in providing an account of intention we should be focusing on the role of intention in planning.

Bratman's Planning Theory of Intention can be contrasted with what appears to be a rather different approach to intention – a strand of thinking about intention that is most associated with Anscombe and which has been championed more recently by David Velleman. According to this approach, the nature of intention is best illuminated by focusing on its epistemological role – in particular, the role intention plays in providing us with a distinctive form of knowledge of what we are doing when we perform our intended actions. The suggestion is that our intentions can embody a form of self-knowledge – practical knowledge of our intended actions. Intended action is behaviour that realises the agent's knowledge of it – practical knowledge of the action that is embodied in the agent's intention.

On the face of it, these two approaches appear to be in tension with one another. According to the one approach the suggestion is that we should focus on the role of intention in planning *future* actions, according to the other we should focus on the role of intention in providing us with a distinctive form of knowledge of what we are doing when we are *performing* our intended actions. The proponent of the former approach claims that future-directed, prior intention should be thought of as paradigmatic, whereas the proponent of the latter approach claims that this will result in a distorted view of the nature of intention. In what follows I will be suggesting that the right account of intention will need to appeal to aspects of both approaches, and that some of the apparent tensions between them can be reconciled by a notion that

connects them – a notion of self-governance. I shall start by briefly outlining some of the key elements of Bratman's Planning Theory of Intention before going on to compare it with Velleman's Epistemic Account.

### **Bratman's Planning Theory of Intention**

Bratman defends the view that intention is a mental state distinct from belief and desire and that our capacity to form intentions is supplementary to our capacity for belief-desire motivation. Our capacity to form intentions allows us to deliberate in advance, when time and resources are more plentiful, and arrive at plans for future conduct. This capacity also allows for inter-personal and intra-personal coordination. Our future-directed plans help us to coordinate our future actions by committing us to future actions that we, and others, can rely on.

In outlining the functional role of intention Bratman suggests that this mental state has both volitional and reason-centred dimensions. Intention, he argues, is conduct-controlling, in contrast to desire, which is conduct-influencing. Intention settles an issue and resists re-consideration of it. It involves dispositions toward further practical reasoning: a disposition to reason about means to intended ends, and a disposition not to form incompatible intentions. On Bratman's view, the coordinating role of intention requires that my intention to  $\Phi$  normally support my expectation that I will  $\Phi$ . This coordinating role requires that one's plans should be internally consistent, consistent with one's beliefs about the world, and means-end coherent.<sup>1</sup> However, Bratman denies that having an intention to  $\Phi$  entails that one believes that one will  $\Phi$ , for he thinks it is possible for there to be cases in which one intends to  $\Phi$  although one is agnostic about whether one will succeed in  $\Phi$ -ing, and indeed he thinks there can be cases in which one intends to  $\Phi$  despite the fact that one is agnostic about whether one will even try to  $\Phi$ . According to Bratman, although knowledge of one's intention to  $\Phi$  can usually *support* the belief that one will  $\Phi$ , it does not require it. In this he is in disagreement with Velleman, who argues that intention is itself a distinctive kind of belief.

## Velleman's Epistemic Theory of Intention

Anscombe is probably the philosopher who is most associated with the idea we have a distinctive form of knowledge of our own actions that is embodied in our intentions. Anscombe argued that expressions of intention, such as “I am going to take a walk”, express the speaker’s knowledge of what he is doing or will do, which is “known by being the content of [his] intention” (1963, p. 53). According to Anscombe, what distinguishes the knowledge embodied in intention from other sorts of knowledge is that it is ‘practical’ knowledge, in the sense that it causes, rather than being caused by, the facts that make it true (p. 87).<sup>ii</sup>

Velleman describes his theory of agency as a variation of this theme of Anscombe’s. According to Velleman most intentions consist in self-fulfilling beliefs about one’s actions. Velleman claims that when a subject decides to act, the intention to act she thereby acquires has the same direction of fit as belief and not desire. Deciding to act does not simply consist in regarding something as *to be* arranged, for once a subject has chosen to act, a question has been settled in her mind. As far as she is concerned, she is going to act. Things have been arranged. The proposition that she is going to act is represented by her as true.

Velleman argues that although intention has the same direction of fit as ordinary belief, it differs in its direction of guidance. According to Velleman, an attitude’s direction of guidance “consists in whether the attitude causes or is caused by what it represents” (2000, p. 25). There can be two ways of accepting something as true: “accepting so as to reflect the truth, and accepting so as to create the truth” (*ibid*, p.194). And so an attitude’s direction of guidance can be either “passive”, or “directive”. Ordinary belief has a passive direction of guidance, whereas intention is directive.

When you decide to do something the intention you thereby acquire is a self-fulfilling expectation about what you’re going to do. Your deciding to  $\Phi$  moves you to  $\Phi$ , and so the expectation that you will  $\Phi$  is self-fulfilling, according to Velleman, because knowing what you are doing is a sub-agential aim regulating your behaviour. You

may have a number of motives for doing various things on a particular occasion, but as your behaviour is guided by this sub-agential aim of knowing what you are doing, in the standard case you will be inhibited from doing anything that you do not expect you are going to do, and so your expectation that you are going to do one of them - your choosing to  $\Phi$  - will result in an additional inclination to  $\Phi$  rather than anything else.

This account of intention is supposed to capture the intuitive thought that when we act intentionally our knowledge of what we are doing is not usually a discovery, “because doing it was our idea to begin with.” The belief is not prompted by prior evidence, because your choosing to  $\Phi$  is what makes it true that you  $\Phi$  rather than anything else, and so, as Velleman puts it, “Although the agent’s expectation of acting is a conclusion to which he jumps before the evidence is complete, he jumps with the assurance that the conclusion will achieve verity even as he lands”. The belief is not entirely unconstrained by evidence, because although directive belief is not prompted by prior evidence, it can be defeated by evidence.

### **Velleman’s Critique of Bratman’s Account**

For Velleman, in contrast to Bratman, the motivational role of intention is not to be explained in terms of its coordinating role in future-directed plans. It is, rather, simply to be explained in terms of the standing desire an agent has to understand what he does: “the best way for him to satisfy this desire is not to do anything until he is prepared to understand it, and then to do that which he is prepared to understand”. In contrast to Bratman, Velleman suggests that our capacity to form intentions “emerges as a byproduct of curiosity plus self-awareness, which are fundamental endowments of human nature, designed for purposes far more general than scheduling deliberation and facilitating coordination”<sup>iii</sup>.

A problem that Velleman raises for Bratman’s account is the frequency with which we form immediate intentions to act now, without any prior planning. The difficulty for Bratman’s account, as he sees it, is to explain why we bother to form such intentions when there is no opportunity for them to play the functional roles that

Bratman specifies. Furthermore, Velleman argues that if, *contra* Bratman, my intention to  $\Phi$  doesn't involve the belief that I will  $\Phi$ , then intention won't, in any case, be able to play the coordinating roles that Bratman identifies, and it won't be subject to the rational requirements that Bratman specifies. Briefly, the thought here is that if I don't believe that I'll do what I intend, then my intention to  $\Phi$  is no longer sufficient as a basis for planning on the assumption that I will  $\Phi$ , and what is more if I don't believe that I'll do what I intend, it becomes unclear why I should limit myself to intentions that can jointly be satisfied in the world as I believe it to be, given that I needn't believe that I will satisfy all of them.

Velleman poses here what I think is an important challenge for Bratman's account of intention, and Velleman's approach to intention has the potential advantages of making an interesting contribution to the epistemology of action.<sup>iv</sup> However, as it stands, I think the account is not without its problems, and solving these problems, I shall suggest, will require incorporating certain elements of Bratman's approach.

### **Problems with Velleman's Account**

Although Velleman thinks that both intention and ordinary belief have the same direction of fit, he does identify a difference between these mental states, namely the fact that they differ in the directions of guidance. What is more, according to Velleman this difference between an intention and ordinary belief is reflected in the propositional contents of these states. According to Velleman an intention to  $\Phi$  is not only self-fulfilling, it is also self-reflective. That is, an intention to  $\Phi$  is a self-fulfilling expectation that one will  $\Phi$  that *represents itself as such*.

With this in mind, let's now consider the question of how it is that a subject knows whether he intends to  $\Phi$  or whether he, rather, has the ordinary belief that he will  $\Phi$ , for it certainly seems that one can have the latter without having the former. Is this simply a matter of the agent knowing that he has a belief with a certain content – a belief with the content that he will  $\Phi$  because of his belief that he will? The problem with this proposal is that it seems that a subject who has no intention of  $\Phi$ -ing can believe that he will  $\Phi$ , and believe that this belief is likely to be self-fulfilling. The

subject who knows this about himself doesn't mistakenly think that he intends to  $\Phi$ . Consider, for example, Parfit's example of the reluctant insomniac who believes both that he will stay awake all night, and that this will happen because he believes it.<sup>v</sup>

Here one wants to say on Velleman's behalf that the subject who intends to  $\Phi$  not only regards as true the proposition 'I will  $\Phi$ ', but thinks that his so regarding it is due to his *intending* to  $\Phi$ , rather than his having the belief that he will  $\Phi$ . But the problem is that this response isn't available to Velleman, as it presupposes a distinction between intention and belief that hasn't yet been identified. A related problem with Velleman's account is that he thinks that the belief that we form when we intend to do something is grounded in general evidence we have concerning our ability to make things happen by forming intentions to do them. The problem this raises for Velleman is the following: How are we to characterise this capacity that an agent takes himself to have? If the agent takes himself to have the general capacity to make things happen by simply forming a *belief* that they will happen, then it seems that the agent takes himself to be omnipotent about the future. If, on the other hand, the agent takes himself to have the capacity to make a limited range of things happen by simply forming the belief that *those* things will happen, then we seem to be denying the possibility of the agent taking himself to be making a prediction that that kind of thing will happen without forming the intention to do it. Here, again, on Velleman's behalf one wants to say that the agent takes himself to have the capacity to make things happen, not by believing that they will happen, but by *intending to do them*, but again, this response isn't available to Velleman as it presupposes an as yet unidentified distinction between belief and intention.

In summary then, a major concern that Velleman raises for Bratman's account is the suggestion that the mental state of prior intention to  $\Phi$  cannot perform the functional, coordinating role that Bratman identifies, and it won't be subject to the rational requirements that Bratman specifies, unless the subject who intends to  $\Phi$  thereby believes that he will  $\Phi$ . However, a difficulty with Velleman's alternative proposal is that it doesn't seem to accommodate adequately the differences between ordinary belief and intention.<sup>vi</sup> My suggestion will be that some kind of synthesis of elements of each approach will help to meet the different challenges that each account faces.

## **Intention, Commitment and Self-Governance**

Both Bratman and Velleman seem to think of intention as involving some kind of commitment on the part of the agent. For Bratman, when an agent decides to  $\Phi$ , he is committed to performing that action, and this commitment serves to control, or constrain, what the agent does in various ways. For Velleman, when a subject decides to  $\Phi$ , a question has been settled in his mind. As far as the agent is now concerned, he is going to  $\Phi$ . The proposition that 'I am going to  $\Phi$ ' becomes one of the subject's commitments. However, it seems that there are two distinct senses of commitment in play here. For Velleman the subject's commitment is some belief he has, whereas for Bratman the way in which the notion of commitment is being used seems closer to the notion of commitment as obligation, as in the excuse 'I have a prior commitment'. In Bratman's account the idea seems to be that the agent who makes plans for future action makes decisions about what to do, thereby forms future-directed intentions to do those things, and in making these decisions the agent *commits himself* to doing those things. But how are we to understand this notion of committing oneself to doing something? Committing oneself to doing something is best thought of, I suggest, as a form of self-binding, for when one commits oneself to doing something one imposes certain constraints on one's own future conduct. I want now to pause briefly to consider the question of what is involved in the notion of self-binding in general.

There are various ways in which one can bind oneself and thereby impose constraints on one's own future conduct – for example, doing things to reduce opportunities for future  $\Phi$ -ing, or imposing significant costs on one's  $\Phi$ -ing. But why should an agent feel the need to bind himself in any way? Why should he impose any constraints on his own behavior? Why shouldn't we just do whatever we happen to feel like doing at the time? Elster has argued that since we are imperfectly rational there are sometimes benefits from having fewer opportunities rather than more, and he has suggested that agents bind themselves (they 'pre-commit' themselves) in order to protect themselves against passion, preference change, and varieties of inconsistent time preference. A presupposition of this kind of suggestion is that the agents in question are able to regard themselves as temporally extended agents, in the sense that they have the

capacity to think of themselves as having different impulses and inclinations at different times.<sup>vii</sup> As Velleman puts it, in the context of a discussion of Kantian ethics, when you have such a capacity you are able to attain “a perspective that transcends that of your current momentary self.”<sup>viii</sup> One has the capacity to think of oneself as one and the same agent who once wanted to  $\Phi$ , but no longer does, or as one and the same agent who now wants to  $\Phi$ , and who later won’t want to  $\Phi$ , or as one and the same agent who now wants to  $\Phi$  and who may later regret  $\Phi$ -ing.

Velleman suggests that when you have the capacity to have his kind of perspective on yourself, you feel an impulse to find a constant perspective on the question the question of what to do when.

This impulse is unavoidable as soon as the more encompassing vantage point appears. As soon as you glimpse the possibility of attaining a constant perspective from which to reflect on and adjudicate among your shifting preferences, you are drawn toward that perspective, as you would be drawn toward the top of a hill that commanded a terrain through which you had been wandering.<sup>ix</sup>

I take it that the subject who attains this perspective on himself thereby regards himself as the kind of agent who can occupy such a perspective.<sup>x</sup> So for such a subject, the question of what to do becomes the question of what he, as agent who occupies this encompassing perspective, should do, and so such a subject no longer *simply* has reason to do whatever he happens to feel like doing at the time. Such a subject can now have a reason not to be dictated to by his current impulse<sup>xi</sup>, as he regards himself as an agent whose perspective transcends that of his immediate felt impulse to act. In having such a reason, the agent who is capable of attaining this encompassing perspective on himself has the ability to govern his own conduct in a way that the agent who cannot attain such a perspective on himself does not.

So the sort of agent who is capable of taking the kind of measures that Elster considers in order to impose constraints on his own future conduct already has a more basic ability to impose constraints on his own behaviour – an ability more basic than

his ability to *pre*-commit himself. My suggestion is that just such an ability is in play when a subject commits himself to doing something. In particular, I shall suggest that when one commits oneself to doing something, one imposes constraints on one's own future conduct by imposing on oneself an obligation to that thing.

To have the capacity to impose on oneself an obligation to do something is to have the capacity for self-governance.<sup>xii</sup> Metaphorically speaking, there's a sense in which the self-governing agent must simultaneously occupy the role of legislator and legislatee. That is, in order for an agent to be capable of governing himself, he must be capable of both imposing obligations on himself, as legislator, and he must be capable of recognising and acting on those obligations, as the one being legislated to. His authority as self-governing legislator depends upon his own recognition of that authority. In fact it is necessary and sufficient for it. If he doesn't recognise the authority of his own legislations, then he cannot be self-governing, for he will have no authority over himself, but if he does recognise the authority of his own legislations, then he has that authority, and so is self-governing. So all an agent needs to do in order to impose on himself an obligation to do something is to recognize the authority of that self-imposed obligation. But what is it for an agent to recognize the authority of his self-imposed legislations? I suggest that it is just a matter of the agent acting in a way, or being disposed to act in a way, that manifests his recognition of these self-imposed legislations. The self-governing agent takes himself to have authority over himself, and he manifests this stance towards himself in the way that he acts. That is to say, the self-governing agent can impose constraints on himself by simply behaving in a way that manifests his recognition of the constraints that he has imposed on himself.

If deciding to do something is a matter of committing oneself to doing that thing, and committing oneself to doing that thing is a matter of imposing certain constraints on one's own future conduct, what are these self-imposed constraints, and how does one manifest one's recognition of them? When one commits oneself to  $\Phi$ -ing, when one imposes on oneself an obligation to  $\Phi$ , a constraint one imposes on one's future conduct is that of avoiding doing things incompatible with  $\Phi$ -ing as well as doing things that will make it more likely that one will  $\Phi$ . Behaviour that manifests one's

recognition of this self-imposed constraint will include not only attempting to  $\Phi$  when the appropriate time arrives, but also engaging in the mental action of reasoning about the means required for  $\Phi$ -ing, making further decisions about one's future conduct on that basis, and avoiding the imposition of further obligations on oneself the discharging of which would make it impossible for one to  $\Phi$ . Recall that the dispositions to reason about means to intended ends and not to form incompatible intentions are part of Bratman's functional characterisation of the mental state of intention. I am now suggesting that these dispositions are best thought of as dispositions to behave in ways that manifest one's recognition of a self-imposed obligation, and the rational requirements that Bratman claims that intentions are subject to can also be seen in this light.

If the behaviour that manifests one's recognition of the constraints one imposes on oneself when one decides to do something includes reasoning about the means required for  $\Phi$ -ing, and avoiding the imposition of further obligations on oneself the discharging of which would make it impossible for one to  $\Phi$ , then arguably a constraint one has thereby imposed on oneself, one's recognition of which is manifested in one's behaviour (i.e. one's practical deliberation), is to treat 'I will  $\Phi$ ' as an assumption in one's practical reasoning. That is to say, the behaviour involved in manifesting one's recognition of the obligation one imposes on oneself when one decides to  $\Phi$  includes behaving as though 'I will  $\Phi$ ' is true.

When one decides to  $\Phi$ , 'I will  $\Phi$ ' is a background assumption governing one's further planning. So when one decides to  $\Phi$  one regards the proposition that 'I will  $\Phi$ ' as true, not because one has acquired the belief that one will  $\Phi$ , but rather because to so regard the proposition is a constraint one necessarily imposes on oneself in making that decision. If one doesn't assume that one will  $\Phi$  when one is engaged in further practical reasoning, then one does not behave in a way that manifests one's recognition of the self-imposed obligation to  $\Phi$ , which amounts to failing to impose that obligation on oneself. If one decides to  $\Phi$  and then subsequently fails to assume that one will  $\Phi$  when one is engaged in further practical reasoning, then one has either forgotten one's intention to  $\Phi$ , which arguably amounts to no longer having that intention, or one has changed one's mind.<sup>xiii</sup>

So now the suggestion here is that Velleman is right to claim that when a subject decides to  $\Phi$  that subject regards as true the proposition 'I will  $\Phi$ '. However, this is not because he has the ordinary belief that he will  $\Phi$ , but rather because he *assumes* that he will  $\Phi$ . Strictly speaking then, 'I will  $\Phi$ ' is not the content of the subject's intention, it is not the content of a special kind of belief that is identical with the intention, but it is, rather, the content of an assumption that a subject necessarily makes when he makes the decision, and thereby forms the intention, to  $\Phi$ . His assumption that he will  $\Phi$  is not grounded in evidence that he possesses for the truth of that proposition, in the way that a prediction might be, but it is, rather, a constraint he imposes on himself when he decides to  $\Phi$ .<sup>xiv</sup> He imposes this constraint on himself by manifesting his recognition of it, that is, by behaving as though he will  $\Phi$  - e.g. by engaging in the mental action of practical reasoning on the background assumption that he will  $\Phi$ .

One cannot impose on oneself an obligation to do something without knowing that this is what one is doing, for the imposition of an obligation on oneself requires behaving (or being disposed to behave) in a way that manifests one's *recognition* of it. So when one decides to  $\Phi$  one not only behaves as though one will  $\Phi$ , but one also knows why one is doing this. One knows that one is assuming that one will  $\Phi$  because one has decided to  $\Phi$ , and so one knows that one's assumption that one will  $\Phi$  is not grounded in evidence in the way that ordinary belief is. So the subject who imposes this constraint on himself knows why he is behaving as though he will  $\Phi$ , and so there is no possibility of his misidentifying his assumption that he will  $\Phi$  for the ordinary belief (i.e. prediction) that he will  $\Phi$ .

If we think of intention in these terms we appear to be able to resolve the difficulties that Velleman raises for Bratman's account, and also certain difficulties that I identified earlier for Velleman's account. On the former, recall that Velleman objects that if don't believe that I'll do what I intend, then my intention to  $\Phi$  is no longer sufficient as a basis for planning on the assumption that I will  $\Phi$ . On the view now being recommended, it is not possible for a subject to decide to  $\Phi$  without assuming

that he will  $\Phi$  when he engages in further planning. To decide to  $\Phi$  is to impose on oneself an obligation to  $\Phi$ . One cannot impose on oneself this obligation to  $\Phi$  without behaving in a way that manifests one's recognition of it, and one cannot behave in a way that manifests one's recognition of it without assuming that one will  $\Phi$  when one engages in further planning.

On the latter, the problem for Velleman, as I see it, is that because he thinks that the mental state of intention is a propositional attitude, and he is led to the view that the attitude to the propositional content of the state is that of acceptance, the only way he can accommodate the distinction between intention and ordinary belief from the agent's own perspective is by reflecting this difference in the content of the state. However, the most natural way of doing this, namely by including in the content of the state a reference to itself as intention as opposed to belief, isn't available to him, as it presupposes a distinction between these states that his account leaves unexplained.

On the view I am suggesting 'I will  $\Phi$ ' is not the content of the subject's intention, it is, rather, the content of an assumption the subject necessarily makes when he decides to  $\Phi$ . This assumption is not a belief based on evidence, but is, rather, a constraint the subject imposes on his own future conduct when he decides to  $\Phi$ . I have suggested that one cannot impose on oneself an obligation to do something without knowing that this is what one is doing<sup>xv</sup>, for to impose the obligation on oneself is a matter of behaving in a way that manifests one's *recognition* of it. So when one decides to  $\Phi$  one not only behaves as though one will  $\Phi$ , but one also knows why one is doing this. One knows that one is assuming that one will  $\Phi$  because one has decided to  $\Phi$ .<sup>xvi</sup> With this distinction between ordinary belief and intention in play, we can now say that the subject who intends to  $\Phi$  assumes that he will  $\Phi$ , and he also knows that he is making this assumption because he *intends* to  $\Phi$ , and not because he takes himself to have the belief that he will  $\Phi$ . We can now also say that this subject takes himself to have the capacity to make things happen, not by believing that they will happen, but by *intending to do them*.

## **Intention-in-Action and Knowledge of Action**

For this account the question of whether intention can embody a form of practical knowledge now becomes the question of whether one's assumption that one will  $\Phi$ , which one necessarily makes when one decides to  $\Phi$ , can embody a form of practical knowledge. Well let's consider Velleman's claims about the direction of guidance of the state of intention, now understood in terms of the direction of guidance of the assumption one makes when one decides to  $\Phi$ . So far we have the idea that reaching a decision to  $\Phi$  involves binding oneself in various ways. A constraint one imposes on oneself when one reaches a decision to  $\Phi$  is to assume that one will  $\Phi$ . When one imposes on oneself an obligation to do something, one assumes that one will do it. This is part of what is involved in behaving / being disposed to behave in a way that manifests one's recognition of this self-imposed obligation, which in turn, is what is required for imposing this obligation on oneself in the first place.

The self-binding involved in deciding to do something makes it more likely that one will perform the action, for it helps to prevent one from deviating from that course of action under the influence of other desires and inclinations, and what is more, it constrains one in a way that allows for means-end reasoning – reasoning which itself makes the performance of the action more likely. So when one intends to  $\Phi$ , one adds to whatever likelihood there was that one would end up  $\Phi$ -ing that obtained simply in virtue of one's prior motivational states.<sup>xvii</sup> This suggests that when one decides to do something that one is capable of doing, the assumption that one will perform the action that one thereby makes has the direction of guidance that Velleman attributes to intention. The state of intention involves regarding as true a proposition that will be (or is usually) made true because one so regards it. But how might this yield an account of how we know what we *are* doing when we are acting? In order to address this question we need to consider the notion of intention-in-action / present-directed intention and its relation to prior intention.

If I commit myself to  $\Phi$ -ing, then what I commit myself to, the object of my commitment, we might say, is doing that thing – that is, me  $\Phi$ -ing. Me  $\Phi$ -ing is what would discharge my self-imposed obligation to  $\Phi$ . If I commit myself to  $\Phi$ -ing

tomorrow, then the object of my commitment is not 'I will  $\Phi$  tomorrow' but rather, me  $\Phi$ -ing tomorrow. When tomorrow arrives, me  $\Phi$ -ing today is what would discharge my commitment. And here we want to say that in an important sense the object of my commitment has not changed since yesterday, although the temporal relation I stand to it has, and so what I must do in order to manifest my recognition of the obligation I imposed on myself in making that commitment yesterday has changed. For in order to manifest today my recognition of the obligation I imposed on myself yesterday when I decided to  $\Phi$  tomorrow, I have to assume that I will  $\Phi$  today. I have to be disposed to behave as though I will  $\Phi$  today when I engage in practical reasoning.

Central to the notion of a self-imposed obligation to  $\Phi$  is the idea of doing things, and being disposed to do things, that manifest one's recognition of that self-imposed obligation, and one of the things one does that manifests one's recognition of the self-imposed obligation to  $\Phi$  is one's actual  $\Phi$ -ing itself. Me  $\Phi$ -ing today, or attempting to  $\Phi$  today, is itself part of what is involved in manifesting my recognition of the obligation I imposed on myself yesterday when I decided to  $\Phi$  tomorrow, and it is what is required for discharging my self imposed obligation to  $\Phi$ . However, one's  $\Phi$ -ing can only count as a case of discharging one's commitment to  $\Phi$  if it is a manifestation of one's recognition of it, and it can only be a manifestation of one's recognition of it if one's commitment to  $\Phi$ -ing obtains while the action is being performed. For if, when one performs the action, one is no longer committed to doing it, if one once was committed to doing it but no longer is, then the action cannot be regarded as one's fulfilling / discharging that commitment. The commitment to  $\Phi$  must obtain while one is  $\Phi$ -ing if the action is to be regarded as a case in which one is discharging one's commitment to  $\Phi$ , and so embedded in the notion of the kind of prior intention that makes planning for the future possible is a notion of intention-in-action – a state that obtains during the performance of the action.

What then distinguishes intention-in-action from prior intention? The intention-in-action obtains only if either there occurs the action that is its object, or there occurs an action that is an attempt to perform the action that is its object. If no such action occurs then the state does not obtain. Either the prior intention, and not the intention-

in-action, obtains (perhaps the subject has failed to keep track of the time) or neither the intention-in-action, nor the prior intention obtains (perhaps the subject has changed his mind, or simply forgotten what he decided to do). In the case of intention-in-action we have a state whose obtaining requires the occurrence of some event – we have a state that only obtains while some occurrence *manifests* it, where that occurrence is its object or an attempt to perform the action that is its object.

In summary, the intention-in-action has the same object as the prior intention – i.e. me  $\Phi$ -ing. The temporal relation I stand in to that object changes as the time for action approaches, but not the object itself. The propositional content I treat as true, the content of the assumption I make in having made the decision to  $\Phi$ , changes as the time for the action approaches. If the content of the assumption does not change then I am failing to recognise and *keep track* of the obligation I imposed on myself yesterday when I decided to  $\Phi$  tomorrow.

So what happens to the subject's future-tense assumption that he *will*  $\Phi$  (in particular his assumption that he *will start*  $\Phi$ -ing) once the action commences?<sup>xviii</sup> Assume that we have a case in which although the subject starts  $\Phi$ -ing, from his perspective he has yet to start fulfilling his self-imposed obligation to  $\Phi$ . In such a case although the subject is actually  $\Phi$ -ing, his  $\Phi$ -ing cannot be a manifestation of his *recognition* of his self-imposed obligation to  $\Phi$ . So if the subject's  $\Phi$ -ing is to be a manifestation of his recognition of his self-imposed obligation to  $\Phi$ , then the subject cannot be assuming that he has yet to start fulfilling this self-imposed obligation. When the subject's  $\Phi$ -ing manifests his recognition of his self-imposed obligation to  $\Phi$  - when his  $\Phi$ -ing is an instance of discharging his self-imposed obligation to  $\Phi$  - then the subject no longer assumes that he *will* start  $\Phi$ -ing – the subject is no longer disposed to behave as if he will start  $\Phi$ -ing. But why does the subject stop assuming that he will start  $\Phi$ -ing? It isn't because the subject is no longer committed to  $\Phi$ -ing. So is it because the subject *discovers* that he has started  $\Phi$ -ing?

In a case in which one discovers that one is doing what one obliged oneself to do, it seems that one's action cannot be a manifestation of one's own recognition of that

self-imposed obligation. For when one's behaviour manifests one's recognition of one's self-imposed obligation to do something one thereby knows why one is doing what one is doing, in so far as one is acting *in recognition* of that self-imposed obligation. So returning to the question of why a subject stops assuming that he *will* start  $\Phi$ -ing when his  $\Phi$ -ing manifests his recognition of his self-imposed obligation to  $\Phi$ : if it isn't because the subject is no longer committed to  $\Phi$ -ing, and it isn't because he discovers that he is  $\Phi$ -ing, it must be because he is assuming that he is  $\Phi$ -ing, or at least assuming that he is attempting to  $\Phi$ .<sup>xix</sup>

If the subject's  $\Phi$ -ing is to be a manifestation of his recognition of the obligation he imposed on himself when made the decision to  $\Phi$ , then the subject must assume that he is  $\Phi$ -ing, or attempting to  $\Phi$ , when he is  $\Phi$ -ing. What we have now arrived at is a view according to which when one's  $\Phi$ -ing is a manifestation of one's intention to  $\Phi$  one is in a state of intention-in-action, which is a state that can obtain only if (a) one is either  $\Phi$ -ing or attempting to  $\Phi$ , and (b) one is assuming that one is  $\Phi$ -ing (or assuming that one is attempting to  $\Phi$ ), where this assumption is not grounded in evidence one has about what one is doing, in so far as it is not a discovery. So there are three components to keep in mind here: the state of intention-in-action, the action, and the assumption that has an action as object – an assumption one necessarily makes when one is in the state of intention-in-action. With this picture in mind, what conclusions can we draw about the epistemic role of the intention-in-action? In particular, what conclusions can we draw about the epistemic status of the assumption involved?

The view I've outlined allows for the following possibility. One intends to  $\Phi$ , one attempts but fails to  $\Phi$ , where this failed attempt at  $\Phi$ -ing is a manifestation of one's recognition of one's self-imposed obligation to  $\Phi$ , and during the attempt one assumes that one is  $\Phi$ -ing. So the assumption one makes when one is in the state of intention-in-action can certainly fail to be knowledge, for it can be false. However, it doesn't follow from this that such assumptions can never embody a form of knowledge. When one tries and fails to  $\Phi$ , where this attempt at  $\Phi$ -ing manifests one's recognition of one's self-imposed obligation to  $\Phi$ , one tries and fails to fulfil

one's self-imposed obligation to  $\Phi$ , and one thereby fails to make true the assumption that one is  $\Phi$ -ing. But suppose that an attempt at  $\Phi$ -ing that manifests one's self-imposed obligation to  $\Phi$  is successful, so one's  $\Phi$ -ing fulfils one's self-imposed obligation to  $\Phi$ . And suppose too that one's success is not accidental. One has the reliable ability to fulfil the self-imposed obligation to  $\Phi$  and one has exercised that reliable ability on this occasion. Then it follows that one has a reliable ability to make true one's assumption that one is  $\Phi$ -ing that one makes when fulfilling the self-imposed obligation to  $\Phi$ , and one has exercised that ability on this occasion. The concurrence of one's  $\Phi$ -ing and the truth of one's assumption that one is  $\Phi$ -ing is non-accidental. It is a reflection of one's reliable ability to fulfil a self-imposed obligation to  $\Phi$ . Given that the concurrence of one's  $\Phi$ -ing and the truth of one's assumption that one is  $\Phi$ -ing is non-accidental, a case can be made for thinking that the assumption thereby embodies a form of knowledge of the action it concerns. One's  $\Phi$ -ing realises one's knowledge of it.<sup>xx</sup>

The assumption that one is  $\Phi$ -ing that one makes when one's  $\Phi$ -ing manifests one's recognition of one's self-imposed obligation to  $\Phi$  is not grounded in evidence one possesses, in so far as one does not *discover* via observation that one is  $\Phi$ -ing.<sup>xxi</sup> However, this assumption, though not grounded in such evidence, can be defeated by it. So when one possesses evidence that undermines one's assumption that one is  $\Phi$ -ing, one ought not assume that one is  $\Phi$ -ing. But recall that the state of intention-in-action cannot obtain without the occurrence of an event that is at least an attempt to perform the action, and what one is doing cannot be a manifestation of one's recognition of one's self-imposed obligation to  $\Phi$  unless one at least assumes that one is attempting to  $\Phi$ . So even if one possesses evidence that undermines an assumption that one is  $\Phi$ -ing, one is still entitled to the assumption that one is attempting to  $\Phi$ . It is always the case that the epistemic residue of this defeating evidence is knowledge that one is attempting to  $\Phi$ .<sup>xxii</sup>

I have suggested that a decision to do something is a form of self-binding. When one decides to do something one imposes constraints on one's future conduct by imposing on oneself an obligation to do that thing.<sup>xxiii</sup> When one imposes on oneself an

obligation to  $\Phi$ , a constraint one thereby imposes on oneself is that of assuming that one will  $\Phi$  when engaged in further practical deliberation. However, although I introduced the account by focusing, as Bratman does, on the future-directed case, I have argued the implications of the account have much in common with what Velleman says about the epistemic role of intention. For a consequence of this account is that when one's  $\Phi$ -ing manifests one's recognition of one's self-imposed obligation to  $\Phi$ , one assumes that one is  $\Phi$ -ing (or attempting to  $\Phi$ ) and this assumption can embody a form of practical knowledge of what one is doing.

Velleman is right to assert that intention can embody a form of practical self-knowledge that is realised by one's actions. However the actions that realise this form of self-knowledge also realise one's self-governance. And central to this notion of self-governance is the subject's ability to attain a perspective on himself as a temporally extended agent – a perspective the attainment of which allows the subject to engage in future-directed planning. So these are also notions that Bratman is right to emphasise in his account of intention. It is the notion of self-governance that is central to an account of intention, and it is this notion that provides the crucial connection between the mental actions of practical deliberation and planning, and the kind of practical self-knowledge that intention can embody and that one's actions can realise.

### **References:**

Anscombe, G.E.M. 1963, *Intention*, 2<sup>nd</sup>. ed., Cornell University Press, Ithaca, NY

Bratman, Michael. 1984, 'Two Faces of Intention', *Philosophical Review* 93: 375-405

Bratman, Michael. 1987, *Intention, Plans, and Practical Reasoning*, Harvard University Press, Cambridge, MA

Bratman, Michael. 1992. 'Practical Reasoning and Acceptance in a Context', *Mind*, Vol. 101, No. 401

Bratman, Michael. 1999, *Faces of Intention: Selected Essays on Intention and Agency*, Cambridge University Press, Cambridge

Bratman, Michael. 2000. 'Reflection, Planning, and Temporally Extended Agency', in *The Philosophical Review* 35-61.

Bratman, Michael. 2006. *Structures of Agency: Essays* (Oxford University Press)

Bratman, Michael. (Forthcoming) 'Intention, Belief, Practical, Theoretical', in Jens Timmerman, John Skorupski, and Simon Robertson, eds., *Spheres of Reason*.

Castañeda, Hector-Neri, 1975, *Thinking and Doing*, Reidel, Dordrecht

Elster, J. 1979. *Ulysses and the Sirens* (Cambridge: Cambridge University Press).

Elster, J. 2000. *Ulysses Unbound*. Cambridge University Press.

Frankfurt, Harry, 1988, *The Importance of What We Care About*, Cambridge University Press, Cambridge

Frankfurt, Harry, 1999, *Volition, Necessity, and Love*, Cambridge, Cambridge University Press

Hampshire, S; Hart, H. L. A. 1958. 'Decision, Intention and Certainty', *Mind*, Vol. 67, No. 265

Harman, Gilbert. 1976. 'Practical Reasoning', *Review of Metaphysics* 79: 431-63

Harman, Gilbert, 1986, *Change in View*, MIT Press, Cambridge, MA

Holton, R. 1999. 'Intention and Weakness of the Will', *Journal of Philosophy* Vol. 96, No. 5

Holton, R. 'Partial Belief and Partial Intentions' (Author's website)

Kant, I., 1785, *The Groundwork of the Metaphysics of Morals*, H.J. Paton (trans.), New York: Harper & Row Publishers, 1964.

Kant, I., 1788, *Critique of Practical Reason*, trans. by L.W. Beck, Indianapolis: Bobbs-Merrill, 1956.

Kenny, A. 1963. *Action, Emotion and Will*. Routledge and Kegan Paul

Korsgaard, C., 2002. Self-Constitution: Action, Identity and Integrity. *The Locke Lectures*

Moran, Richard, 2001, *Authority and Estrangement: An Essay on Self-Knowledge* Princeton University Press, Princeton

Nagel, T. 1979. *The Possibility of Altruism*. Princeton University Press.

Parfit, D. 1984. *Reasons and Persons*. Oxford University Press

Peacocke. C. (1999). *Being Known*. Oxford University Press.

Pickard, H. 2004. 'Knowledge of Action Without Observation' *Proceedings of the Aristotelian Society*

Roessler, J. 2003. 'Intentional Action and Self-Awareness', in J. Roessler & N. Eilan (eds.), *Agency and Self-Awareness*, Oxford University Press

Searle, John, 1983, *Intentionality*, Cambridge University Press, Cambridge

Taylor, C., 1976, 'Responsibility for Self,' in A. Rorty ed. *The Identities of Persons*, Berkeley: University of California Press.

Velleman, David, 1989, *Practical Reflection*, Princeton University Press, Princeton

Velleman, J. D., 1989a, 'Epistemic Freedom' *Pacific Philosophical Quarterly* 70: 73-97.

Velleman, David, 2000, *The Possibility of Practical Reason*, Oxford University Press, Oxford

Velleman, J. D., 2000b "What Happens When Someone Acts?" in Velleman 2000 123-43.

Velleman, J. D., 2000c. 'The Possibility of Practical Reason', in Velleman 2000.

Velleman, J. D., 2005a. 'Willing the Law', in *Self to Self: Selected Essays*. Cambridge University Press

Velleman, J. D., 2005b. 'A Brief Introduction to Kantian Ethics', in *Self to Self: Selected Essays*. Cambridge University Press

Velleman, J. D., 'What Good is a Will?' (Author's website)

Watson, Gary (1982a). "Free Agency," in Watson ed. *Free Will*. Oxford: Oxford University Press.

Williams, B. 1981. 'Persons, Character and Morality', in *Moral Luck*. Cambridge University Press.

---

<sup>i</sup> According to Bratman, "these demands are rooted in a pragmatic rationale: their satisfaction is normally required for plans to serve well their role in coordinating and controlling conduct" (1987, p. 31-32).

<sup>ii</sup> See also Hampshire and Hart 1958.

<sup>iii</sup> 'What Good is a Will?', p. 25

<sup>iv</sup> As well as providing a novel explanation of the openness each agent sees in his own future. See Velleman 1989a.

<sup>v</sup> Parfit's example is discussed by Harman (1976), and also Bratman (forthcoming).

<sup>vi</sup> Can these difficulties be resolved by supplementing Velleman's account with some further characterisation of the distinction between ordinary belief and intention? Harman (1976), for example, has suggested that the difference between belief and intention should be understood in terms of the origin of these attitudes – i.e. in terms of whether theoretical or practical reasoning results in the acquisition of the attitude. However, there appear to be two potential problems with this suggestion as

---

it stands, analogous to those that apply to Velleman's account. Unless we are given some reason for thinking that a subject cannot be wrong about the causal origin of the mental state he is in we would then seem to allow for the potential possibility of misidentifying one's prediction as an intention, and vice-versa. (Compare here the passage in Anscombe in which she briefly discusses the impossibility of misidentifying one's intention – 1963, section 3). Furthermore, the suggestion appears to presuppose that we should be able to distinguish between theoretical and practical reasoning independently of the distinction between decision / intention and ordinary belief, and it is not obvious how this is to be done.

<sup>vii</sup> Note this capacity isn't simply to be understood as the capacity to do something now with the aim of getting something later. Compare here Elster's discussion (in *Ulysses and the Sirens*) of the difference between the capacity to employ indirect strategies and to enact behaviour that delays the attainment of a goal on the one hand, and the *generalised* capacity for global maximization on the other: "The characteristic of man is not a programmed ability to use indirect strategies or adopt waiting behaviour in specific situations, but rather a *generalized* capacity for global maximization that applies even to qualitatively new situations." (p. 16)

<sup>viii</sup> Velleman 2005b.

<sup>ix</sup> Velleman 2005b, p. 23. I don't think that this in itself entails that one will thereby be drawn to make decisions concerning one's actions in a temporally neutral way. That is, the decisions one makes from such a perspective may still reflect a bias towards the present, a bias towards the near-in-time, and in some cases a bias towards the future. (See the discussion of these possibilities in Parfit, *Reasons and Persons*, Part II). A question that then arises is whether, given the subject's ability to attain such an encompassing perspective, these forms of time preference necessarily involve a form of irrationality. For the claim that they don't, see Williams, 'Persons, Character and Morality'. See also the discussion in Parfit, and Nagel, *The Possibility of Altruism*.

<sup>x</sup> It might be argued that the subject who shows a preference for the present doesn't always fail to regard himself as an agent who occupies this encompassing perspective, for his current decision might in part be based on the fact that he trusts himself to make decisions from this encompassing perspective at a later time, in which case he thereby implicitly acknowledges his own ability to make decisions from this encompassing perspective. But whether the agent would be irrational in regarding things in this way is a further matter. For the suggestions as to how such a view can lead to inconsistency see Elster *Ulysses and the Sirens*, and *Ulysses Unbound*.)

<sup>xi</sup> Where this current impulse might include the impulse to do something now with the aim of getting something later. See earlier footnote.

<sup>xii</sup> The notion of self-governance that I am using here differs from the one that Velleman discusses in *Practical Reflection* and 'The Possibility of Practical Reason'. In *Practical Reflection* Velleman naturalises the Kantian idea of the will as a capacity to act in accordance with one's own idea of the law. He suggests that the notion of law here should not be understood not as a normative law, but as a causal one. So the will is the capacity to act in accordance with one's own idea of the causal laws that govern one's actions. The will is "a faculty for writing part of our own constitution in the form of a theory of how we are constituted". The idea of self-governance as self-imposed obligation doesn't appear to be any part of this account. Similarly, in 'The Possibility of Practical Reason', where Velleman argues that the constitutive aim of action is autonomy (a view he later rescinds, and replaces with the idea that the aim of action is self-understanding), the notion of autonomy he has in mind is that of a controlling consciousness of one's behaviour – a 'guided awareness' of one's action. Again, the notion of self-imposed obligation does not appear to play any part in this account of self-governance. On a related point, in *Practical Reflection*, Velleman mentions that he considers his theory to be a variation of the view of intentions as commands. (See Kenny 1963 and Casteneda 1975). He doesn't identify intentions with commands, but he does identify them with expectations whose utterance would convey commands. Velleman defends the idea that a single utterance can both be the expression of an expectation and a command, and his suggestion seems to be that when a subject forms an intention to do something, there's a sense in which the reflective self utters an expression of an expectation and thereby makes a command to the acting self. However, it is unclear to me why we should think that under Velleman's account any *command* has been issued or obeyed. For it seems that under Velleman's naturalised account, the fact that one acts on one's expectation to act is simply explained in terms of one's natural disposition to do what one expects in order to understand what one is doing, and not because there is any sense in which one is obeying a command. One possible response here might be to take the sort of Sartrean line that Charles Taylor develops, and claim that the *manner* in which one understands oneself affects the constraints one thereby imposes on one's own conduct. As Elster puts it, the idea might be that self-understanding is part discovery, part description and part creation.

---

(See also the discussion in Moran, *Authority and Estrangement*). However, it still does not seem clear to me that this would amount to anything like one's obedience to a command that is directed to oneself.

<sup>xiii</sup> If forgetting that one intends to  $\Phi$  can amount to a way of losing the intention to  $\Phi$  without changing one's mind, then this may be further evidence that intention involves something belief-like. Contrast here the case of desire, a state whose persistence doesn't depend in the same way on one's memory.

<sup>xiv</sup> When one decides to  $\Phi$ , a constraint one thereby imposes on oneself is that of assuming that one will  $\Phi$ . One may then make various other assumptions about other things that are going to happen that are grounded in the assumption that one will  $\Phi$ , but that does not in itself entail that one is thereby intending to do those things. One intends to do those things (in the relevant sense) only if one has decided to do those things, one has decided to do those things only if one has imposed on oneself an obligation to do those things, and if one has imposed on oneself an obligation to do those things, then the assumption that one is going to do those things is not *simply* grounded in further assumptions or beliefs one has. Compare here Hampshire and Hart 1958: "If the action is what he intended to do, then the agent must know what he was doing in some sense which would differentiate his shooting at the bird from other non-accidental actions performed at the same time, such as making the cartridge explode."

<sup>xv</sup> Although one can do something, which, as it happens, puts oneself under some obligation without realising that this is what one has done. E.g. cases in which the act of  $\Phi$ -ing puts one under some legal/moral obligation, but one is ignorant of this fact, or one doesn't realise one is  $\Phi$ -ing.

<sup>xvi</sup> To say that one cannot misidentify one's intention as a belief and to say that one cannot intend to do something without knowing that this is what one is doing is of course not to say that one cannot have a false belief that one intends to do something. (Bratman argues that it is possible for an agent to believe that he intends to  $\Phi$ , when in fact he doesn't intend to  $\Phi$ . See Bratman's 'Intention and Means-End Reasoning'). One might, for example, think one has decided to do something because one took oneself to be deciding to do something, but where one didn't actually thereby acquire a state that satisfies all the relevant manifestation conditions for having such an intention, just as one might think one believes something when one doesn't, because one takes oneself to be judging something, but one fails to acquire a state that satisfies all the relevant manifestation conditions for having such a belief. See the discussion of this possibility with respect to belief in Peacocke, *Being Known*, and Moran, *Authority and Estrangement*.

<sup>xvii</sup> So there is an important difference between the assumption that one will  $\Phi$  that one makes when one decides to  $\Phi$  and a prediction one makes that one will  $\Phi$  which one bases on one's belief about one's desire to  $\Phi$ .

<sup>xviii</sup> An issue I think is worth exploring, but which I don't have the space to discuss here, concerns what should be said about the expression of an intention-in-action. In particular, it is worth comparing the differences between the expression of intention-in-action where the intended action is an activity, and the expression of an intention-in-action where the intended action is an accomplishment, and the way in which this has implications for the tense of expression of intention-in-action, as well as the question of the kind of knowledge one can have of an accomplishment that does not reach its terminus.

<sup>xix</sup> If during the action the subject merely assumes that he is attempting to  $\Phi$ , rather than assuming that he is  $\Phi$ -ing, presumably we will want some explanation of why the subject's prior assumption that he will  $\Phi$  becomes the more tentative assumption that he is attempting to  $\Phi$  as the action commences. So if the subject merely assumes that he is attempting to  $\Phi$  when his  $\Phi$ -ing manifests his self-imposed obligation to  $\Phi$ , then this suggests that either the subject's prior decision was to attempt to  $\Phi$ , as discussed earlier, or the subject has, in the meantime, gained some further evidence concerning his chances of successfully  $\Phi$ -ing. This latter point is one I shall say more about shortly.

<sup>xx</sup> It is in this sense that when the assumption is false the fault lies with the action and not the assumption, and this is related to the idea of one's action *realising* one's knowledge of it. Compare here Anscombe 1963, and also Hampshire and Hart 1958.

<sup>xxi</sup> I take this to be consistent with the claim that perceptual experience in action can be a source of knowledge of what one is doing. See Johannes Roessler 'Intentional Action and Self-Awareness', and Hanna Pickard, 'Knowledge of Action Without Observation'.

<sup>xxii</sup> Note also that in the case of bodily action one's belief about what one is doing can potentially be objectively undermined by evidence one does not possess – e.g. by perceptual evidence that others possess – or perceptual evidence that others could possess. This brings out an interesting potential contrast with the case of certain kinds of mental actions – or at least certain forms of knowledge one can have of such actions. What this account of knowledge of action brings out is that the epistemic difference between knowledge of mental action and knowledge of bodily action may have nothing to

---

do with a special *source* of knowledge of our actions that we have in the case of mental action. The difference may lie in the kind of counterevidence there can be, and that can be available to others and us, in the case of bodily action. The picture of our knowledge of our actions that invokes the notion of an assumption about what we are doing (where this involves the imposition of some kind of self-imposed constraint) suggests that the kind of counterevidence that is available in the case of bodily action, in comparison with the case of certain kinds of mental action, can make an epistemic difference. In particular, it may make knowledge of one's bodily actions susceptible to a form of sceptical attack in a way that one's knowledge of certain mental actions may not be.

<sup>xxiii</sup> In forming the immediate decision to now  $\Phi$ , without any prior deliberation and planning, the agent is still exercising his capacity for self-governance. We are still to regard this decision to  $\Phi$  as the self-imposed obligation to  $\Phi$ , as his  $\Phi$ -ing is a manifestation of his recognition of that self-imposed obligation. In which case, I suggest, although we can attribute this form of practical self-knowledge to a subject when he hasn't engaged in any prior deliberation or planning, we cannot attribute such self-knowledge to an agent who isn't capable of engaging in such prior deliberation and planning. Note that no reason has been offered for thinking that there cannot be action without the decision to perform the action, and no reason has been offered for thinking that there cannot be action without that action manifesting the agent's intention to perform it. So no reason has been offered for thinking that there cannot be action without this form of self-knowledge. Note also that no assumption is being made here that this account of the epistemological role of intention exhausts what should be said about the way in which we know what we are doing when performing actions, and I haven't committed myself on the question of whether an action that doesn't involve this form of practical knowledge can be correctly described as 'intentional'.