Inescapability and the Analysis of Agency[1]

Phil Clark

University of Toronto
philip.clark@utoronto.ca

In *How We Get Along*,[2] Velleman pursues a Kantian strategy. The point of this strategy, he says, is "to show that we can account for the objectivity of morality without positing a normative reality of which judgments of right and wrong can be true (120)." To this end, Velleman tries to explain how the demands of practical reason can have a kind of objectivity even if there is no normative reality of which practical reason judgments can be true. The key is to recognize that while all reasons for action rest on motives that are present in those for whom they are reasons, some rest only on a motive that is necessarily shared by all who are subject to reasons. In this paper I argue that Velleman's theory of reasons, even if it is true, fails to support the hypothesis that there is some one motive that is necessarily shared by all those who are subject to reasons for action. To make good on his Kantian strategy, Velleman must establish this hypothesis in some other way.

I proceed by granting Velleman his theory of reasons. In that theory, action constitutively aims at behaving in a way that makes sense, and reasons for action are indicators of what would make for successful action. In granting this for the sake of argument I depart from other critical strategies, for instance that of Kieran Setiya, who thinks Velleman has misidentified the constitutive aim of action, and therefore rejects Velleman's theory of reasons.[3] Setiya does not question whether Velleman's theory of reasons, if true, would support his Kantian strategy, whereas questioning that is precisely my business here. My concern is also not that of David Enoch, who worries that even if Velleman is right about the constitutive aim of action, he can't explain why we should do actions rather than, say, shmactions, which are like actions except for lacking the aim that makes something an action.[4] Like Setiya, Enoch is questioning Velleman's account of reasons for action, whereas again I am just questioning whether that account, if true, would support the Kantian strategy.

To do this I first need to explain Velleman's Kantian strategy. After that I describe his theory of reasons, and finally I address the relation between the two.

1 The Kantian Strategy

A strategy is for doing something. What does Velleman want the Kantian strategy to do? The answer, I think, is that he wants it to resolve an apparent tension within his view. He wants to "explain how morality can be objective." And he does this by first explaining how

---

[2]All citations in the text are to Velleman (2009). For the Kantian strategy, see also Velleman (2000: Chapter 8).

[3]See Setiya (2003: 376).

[4]See Enoch (2006) and (2011).

practical reason can be objective, and then explaining how practical reason supports morality. The problem is that he also accepts claims that can seem to rule out any sort of objectivism about practical reason. He takes these claims from Bernard Williams:

> I agree with Williams's premise that reasons for acting must be able to engage a motive that the agent has or could come to have through sound deliberation; and I do not wish to question the assumption that deliberation can convey him only from motive to motive, so that his current motives determine where he could rationally end up. But I reject Williams's conclusion, that reasons must therefore be geared to something subjective in the agent's psychological make-up (119-20).

Velleman grants that nothing can be a reason for action unless it either engages a current motive of the agent or engages a motive that could be derived from a current motive of the agent. He thus accepts that reasons must be "geared to" something in the agent's psychological make-up; they must be geared to some current motive of the agent. And this seems to spell doom for any thought that the demands of practical reason, and of any morality spun from them, might be objective. For surely an objective demand would be one that was there in the world quite apart from what the agent happens to want. So here is the problem. If practical reason is tethered to current motives of the agent, how can it be objective?

As Velleman notes, Williams's answer is that it can't be objective. For Williams the demands of practical reason are subjective, because they rest on current motives of the agent. But Velleman thinks this is too quick. Indeed he thinks Williams has shown it is too quick, but has not made enough of his own insight. For elsewhere Williams distinguishes two kinds of objectivity. A demand can be objective in the sense of being "woven into the fabric of the world," but it can also be objective in the sense of being "woven into the practical point of view (116)." The former is the sort of objectivity Mackie famously argued is not to be had in ethics, but the latter is the sort Kant famously argued is to be had in ethics. Velleman quotes Williams making this distinction:

> Consider another picture of what it would be for a demand to be 'objectively valid'. It is Kant's own picture. According to this, a demand will be inescapable in the required sense if it is one that a rational agent must accept if he is to be a rational agent. It is, to use one of Kant's favorite metaphors, *self-addressed* by any rational agent. Kant was wrong, in my view, in supposing that the fundamental demands of morality were objective in this sense, but that is not the immediate point, which is that the conception deploys an intelligible and adequate sense of objectivity. It seems to have little to do with those demands being part of the fabric of the world; or at any rate, they will be no more or less so than the demands of logic – which was, of course, part of Kant's point (115).[5]

Velleman's problem was to explain how the demands of practical reason could be objective given that they are tethered to current motives of the agent. His solution is that they can be objective in the sense of being woven into the practical point of view, even if it they are not objective in the sense of being woven into the fabric of the world. Practical reason and morality can exhibit Kantian objectivity even if they do not exhibit the sort of objectivity Mackie considered and rejected.

---

[5]Bernard Williams, "Ethics and the Fabric of the World," in Williams (1995: 172-81)

The Kantian strategy, then, is to explain how the demands of practical reason are woven into the practical point of view. Velleman's explanation is that there is a motive, namely the motive of self-understanding, that must be present in anyone who is susceptible to reasons for action. Unlike other motives, this is not a motive agents just happen to have. It is not contingent, he says, but inescapable. If there are demands of practical reason that are tethered to this motive and this motive only, then these demands will be objective, not the sense that they are there in the world apart from any motives of the agent, but in the sense that they are reasons for anyone who can have reasons at all.

> For suppose that there is a single motive that any reason must engage in order for an agent to act on the basis of it. A creature will need this motive in order to satisfy the prerequisite for being subject to reasons for acting – in order for there to be reasons *for* him, or *applicable to* him – but he will then satisfy that prerequisite with respect to any and all possible reasons. Variance in motivation will no longer entail that similarly situated agents can have different reasons, . . . it will entail only that some creatures but not others can have reasons at all, because only some can be motivated in the relevant way (120).

This is where Velleman parts company with Williams. He agrees that reasons for action must be geared to current motives of the agent, but denies that they must be geared to motives that can vary from agent to agent. Instead they may rest on the motive of self-understanding, which is inescapable. So the Kantian strategy is to make room for objectivity in practical reason by showing that there is a single motive that is necessarily shared by all who are susceptible to reasons for action. Of the motive of self-understanding Velleman writes:

> Without this aim to make us susceptible to reasons, we would be incapable of acting for reasons, would not be agents, and would therefore be exempt from the force of reasons altogether. . . . Hence reasons for acting depend for their influence on a motive, but it is not a motive dependent on our several subjective constitutions; it's a motive that provides our shared constitution as agents. Reasons are therefore objective, and their status as reasons can be established once and for all, by the philosophical analysis of agency. That's the Kantian strategy (146-7).

It is important to distinguish two things Velleman might be saying here. One is that the fact of a necessarily shared motive makes room for objective reasons. The other is that the fact of a necessarily shared motive guarantees that there are objective reasons.

The latter claim faces a serious objection. It could be that although agents necessarily share the motive of self-understanding, what makes for self-understanding always depends on desires the agent just happens to have. In that case there would be no reasons that held for all agents regardless of their subjective constitutions. To show that there are objective reasons, therefore, it is not enough to show that there is a motive that is necessarily shared by all agents. One must go on to show that there are reasons that rest only on that motive.

While Velleman does not explicitly distinguish the two claims, I am going to assume he intends only the former. What he takes to follow from the fact of a necessarily shared motive is that there may be reasons that are objective despite being tethered to current motives, not that there actually are such reasons.[6]

---

[6]A question thus arises. Does Velleman try to show that there are reasons that rest solely on the motive of self-understanding? As I read the book, the early chapters amount to a sustained argument that groups of humans each

Why should we think, though, that there is some one motive that any reason must engage in order for an agent to act on the basis of it? Perhaps, if I am to act for a reason, that reason must engage some motive or other, but why must it engage the same motive every time? When I close the door because there's a cold draft, that reason engages the motive of staying warm, and when I put on my glasses because I can't read the fine print, that reason engages the motive of reading the fine print, but these are different motives. Why think there must be some one motive that is engaged by both reasons? To answer this question, Velleman turns to his theory of practical reasoning, on which he says his version of the Kantian strategy relies (117).

2   Reasons for Action

Velleman builds his account of reasons for action on an account of action, much as he builds his account of reasons for belief on an account of belief. Where belief is concerned, he starts with a distinction between representing something as true and believing it. I can represent it as true that am slaying a dragon, perhaps in a daydream, without believing I am slaying a dragon. What then is the difference between believing something and merely representing it as true? Velleman's answer is that believing Q is representing Q as true with a certain aim, in particular the aim of truth. So his account of belief goes roughly like this:

> **Analysis of Belief)** For S to believe that Q is for S to represent Q as true in an attempt[7] to represent as true whatever really is true as regards whether Q

He then treats reasons for belief as indicators of what would make for successful belief. Thus for instance, someone's wearing a gold band will be a reason for thinking he is married just in case his wearing a gold band is an indicator that representing it as true that he is married would make for success in an attempt to represent as true what really is true as regards whether he is married. More generally, the account of reasons for belief goes something like this:

> **Theory of Reasons for Belief)** For P to be a reason for S to believe that Q (or to believe that not-Q) is for P to be an indicator of what would make for success in an attempt to represent as true whatever really is true as regards whether Q[8]

Velleman models his account of reasons for action on this way of understanding reasons for belief. He begins with an account of action, and then characterizes reasons for action as indicators of what would make for successful action.[9] Moreover, his approach to action parallels his approach to belief. For belief he started with a distinction between representing as

---

of whom succeeds in molding his or her behavior and attitudes to what makes sense will "in the very long run" develop scenarios or protocols – ways of getting along – that have a distinctively moral cast. In this way, practical reason turns out to be "pro-moral." (The argument culminates at 149-51.) I suppose this might be read as an argument that there are reasons that rest only on the motive of self-understanding.

As far as I can see, however, the argument that practical reason is pro-moral could succeed even on the assumption that all reasons rest on contingent desires of agents. Perhaps over time humans with different subjective constitutions, and hence with different reasons, would under pressure of the necessarily shared drive for self-understanding come to have subjective constitutions, and hence reasons, that favor acting morally. But this is consistent with the hypothesis that all reasons rest on contingent desires. It is not clear, then, how Velleman can reach the conclusion that there are reasons that rest solely on the inescapable motive of self-understanding.

[7]The attempt, Velleman says elsewhere, need not be an attempt on S's part. It might just be an attempt on the part of S's cognitive faculties. See his 2000: 184-5.

[8]This is an account of theoretical reasons for belief only. If there can be practical reasons for belief, of the sort envisaged in Pascal's Wager, then one will need a separate account of these.

[9]One might worry that the notion of an indicator is too close to that of a theoretical reason for belief to be useful in an account of theoretical reasons for belief. But even if that is right, the notion of an indicator might be helpful in account of reasons for action. On this point see Setiya (2003: 339-40).

true and believing. For action he starts with a parallel distinction between mere behavior and action. Behavior can fail to be action. Crying, for instance, "can be a completely involuntary outpouring of emotion," in which one "just lets oneself go." This is mere behavior. Alternatively, one can "indulge in a good cry," in which case one's crying is guided by an idea of what one is doing, a "conception of crying." Depending on one's conception of crying, one may sniffle into tissues, or rend one's garments, or shout out lamentations. This is more than mere behavior, for Velleman. It is action, because in it one brings the manifestation of the emotion into accord with a conception of what one is doing (10-11, 26).[10]

Just as belief is more than representation-as-true, then, action is more than behavior. And just as he explains belief as representing-as-true with a certain aim, so Velleman explains action as behaving with a certain aim. To act, he says, is to behave with the aim of behaving in a way that makes sense. So the account of action goes roughly like this:

> **Analysis of Action)** For A's φ-ing to be an action is for A to φ in an attempt to behave in a way that makes sense as regards φ-ing

Thus the bereaved person's behavior – her crying or not, and her crying in one way rather than another – counts as action just in case she cries, or doesn't cry, or cries this way or that, in an attempt to behave in a way that makes sense. If she shakes uncontrollably, and not in an attempt to cry in a way that makes sense, then her shaking is mere behavior, even if it is a manifestation of her grief.

But is Velleman seriously suggesting that when I close the door to stay warm, or put on my glasses to read the fine print, my ultimate goal is something other than staying warm or reading, namely self-understanding? No he is not. As Velleman sees it, we pursue ends like staying warm for their own sake, not for the sake of understanding ourselves. What makes us agents, though, is that we seek to pursue those ends in ways that make sense. In this respect, he says, self-understanding is like efficiency:

> We cannot pursue efficiency alone; we can pursue it only in the course of pursuing other aims, by seeking to pursue them efficiently. And in seeking to pursue them efficiently, we don't pursue them for the sake of efficiency; we pursue them for themselves, albeit with the additional aim of doing so efficiently. So it is with self-understanding (27-8).

Having described action as "behavior aimed at intelligibility, just as belief is acceptance aimed at truth (133)," Velleman goes on to characterize reasons for action as indicators of what would make for success.[11] So the account of reasons for action parallels the account of reasons for belief, roughly as follows:

> **Theory of Reasons for Action)** For P to be a reason for A to φ (or not to φ, or to φ this way or that) is for P to be an indicator of what behavior on A's part would make for success in an attempt to behave in a way that makes sense as regards φ-ing

---

[10]Note that mere behavior, for Velleman, can be more than mere bodily movement. It can be motivated, in this case by grief, and it can even be a case of taking means to a desired end. On the latter point see his 2000: 189-191.

[11]"Reasons," he says, are "considerations that indicate an action to meet a substantive criterion of aptness or correctness (124)." But "having an aim already establishes a criterion of success or failure, which in turn yields a criterion of correctness for whatever can promote or hinder success (136)." Moreover, "action constitutively aims at making sense (146)," and this aim establishes the criterion of correctness for action. To be a reason for or against an action, then, is to be a consideration that indicates of some behavior that it will promote or hinder success in an attempt to behave in a way that makes sense.

Velleman cashes out the relevant notions of understanding and intelligibility in terms of a non-normative reading of the expression "what makes sense." He notes that we are apt to take this expression as meaning "what's appropriate or right or best." In Velleman's view, however, there is another notion of intelligibility from which appropriateness, rightness and goodness derive (27). The expression "what makes sense," read this way, means what is explicable in folk psychological terms. To make sense of someone's behavior in this sense is to trace the behavior to its causes in that person's motives, traits, and other dispositions. So Velleman does not see action as an attempt to behave appropriately, under that description. He sees it as an attempt to behave in a way that "can be understood as caused by [one's] motives, habits, and other characteristics (185)." He first explains reasons for action as indicators of what will make for success in such an attempt, and then explains appropriateness in terms of reasons.

3   Inescapability and Objectivity

As I said, Velleman takes his version of the Kantian strategy to rely on his view of reasons for action. And one can see why he would. The Kantian strategy, recall, was to explain how the demands of practical reason could be objective, by explaining how they are woven into the practical point of view. The explanation was that the demands of practical reason rest on a current motive of the agent – the motive of self-understanding – that is inescapable. The Kantian strategy thus relies on that motive's actually being inescapable. But this is precisely what Velleman's theory of practical reason provides. If Velleman is right that every action necessarily aims at self-understanding, then the motive of self-understanding is inescapable for agents. And this gives the Kantian strategy what it needs.

Or so it seems. In fact I think that while there is a way in which the motive of self-understanding is inescapable, on Velleman's view of practical reason, this does not give the Kantian strategy what it needs.

To see this, we first need to distinguish two ways in which a motive can be inescapable as opposed to contingent. Consider an analogy with another sort of an attempt, not an attempt to represent the truth or to make sense of one's behavior, but an attempt to locate an object.

Suppose we refer to any attempt to locate something as a search. And suppose we coin the term searcher to stand for anyone who ever searches for anything. Now one thing we might want to say about searchers is that anyone who searches for anything must have the motive of locating that thing:

> **Local Search)** For any object the motive of locating that object is inescapable
> for anyone who searches for that object

This is certainly a claim about inescapability. What it says is that the motive of locating a thing is inescapable for those who search for that thing. Thus the motive of locating my keys is inescapable for anyone who searches for my keys, and the motive of locating your dog is inescapable for anyone who searches for your dog. But it doesn't follow that there is any motive that is necessarily present in all who search for anything. Certainly there can be searchers who don't search for my keys, and searchers who don't search for your dog. And as a general matter there may be nothing that all searchers must try to locate. If so, then we have as yet no reason to think there is any motive that all searchers must share. It may be that each search is propelled by a motive that is contingent for searchers. Each of those motives will be inescapable for those who search for that thing, but contingent for searchers.

If this is right then we need to distinguish the claim above from the claim that there is a motive that is inescapable for searchers. It is one thing to say the motive of locating an object

is inescapable for those who search for that object, and quite another to say there is a motive that all searchers must share. We can put the latter claim as follows:

> **Global Search)** There is a motive such that, necessarily, anyone who searches for anything has that motive.

To pursue the analogy, we now need to consider a specific version of this claim, namely the suggestion that the motive of locating objects is inescapable for searchers. This is analogous to saying that the motive of truth is inescapable for believers, and that the motive of self-understanding is inescapable for agents. To say these things is to say that there is a motive that is necessarily shared by believers or agents, and to say something about what that motive is. But what would such a motive be like? Here too the analogy may help.

What would it mean to say there is a motive that is necessarily shared by all searchers, namely the motive of locating objects? In fact it is not altogether clear what this could mean. This necessarily shared motive of locating objects cannot be the motive of locating my keys, or your dog, because these are motives a searcher could lack. They are contingent for searchers, not inescapable. And even if it were impossible to search for anything without searching for, say, the elixir of life, the motive that was necessarily shared would be the motive of locating the elixir of life, not the motive of locating objects.

A more promising picture is available, however. For suppose that although there is no particular object that I want to locate, I still want there to be an object that I locate. This desire could move me to pick some item, say the world's largest frying pan, and search for it, as a way of bringing it about that I locate something. In that case the motive of locating the world's largest frying pan is derived from the motive of locating objects. One might also think of egg hunts here, or of the modern game of geocaching.[12]

Plainly this is a non-standard way of coming to search for something. The more usual way is to search for things on the basis of a need or desire to locate those things, rather than from any general desire that there be things that you locate. But this example does give us a clear picture of what the motive of locating objects could be.

Now suppose we understand the motive of locating objects in this way. What then becomes of the idea that the motive of locating objects is inescapable for searchers? There are two things to note here. The first is that this idea does not follow from Local Search. The second is that the idea is not very plausible. I'll take these points in turn.

The claim that the motive of locating objects is inescapable for searchers entails the claim that there is a motive that is inescapable for searchers, and adds a specification of what that motive is. So we can formulate it as follows:

> **Global Search 0.1)** There is a motive such that, necessarily, anyone who searches for anything has that motive, and it is the motive of locating objects.

And we can note that Global Search 0.1 entails Global Search.

We have already seen, though, that Global Search does not follow from Local Search. This was because Local Search could be true even if there was nothing every searcher had to search for. And this means Global Search 0.1 doesn't follow from Local Search either. For suppose it did. Then because Global Search 0.1 entails Global Search, Global search would follow from Local Search, and that gives a contradiction.

---

[12]Geocaching is the recreational activity of hunting for and finding a hidden object by means of GPS coordinates posted on a Web site. My authority is the *New Oxford American Dictionary, 2nd Edition*.

So Local Search does not entail Global Search 0.1. Or to put it in English, it doesn't follow, from the claim that the motive of locating an object is inescapable for those who search for that object, that the motive of locating objects, as currently construed, is inescapable for searchers.

But beyond this logical point, the two claims also differ in plausibility. Local Search is almost trivially true, whereas Global Search 0.1 is highly questionable.

As we've seen, the standard way of coming to search for something does not require the motive of locating objects, so understood. Normally, one just sees a need or conceives a desire to locate the object itself. On the face of it, then, there could be searcher who always came to search for things in this way, and never came to search for things in the non-standard way. And if that is so, then it is quite mysterious why there could not be a searcher who lacked the motive of locating objects. It may be common in humans to find some entertainment value in the sheer activity of locating things, but the claim in question is much stronger than that. The claim is that searching for anything entails the presence of that motive, so that there simply could not be a being that searched for things but lacked that sort of general attachment to the activity of locating things. That is what I am saying is implausible.

It is far from obvious, then, that the motive of locating objects is inescapable for searchers. But now what of the claim that the motive of locating an object is inescapable for those who search for that object? This seems a fair candidate for a conceptual truth. We simply stipulated that searching for something was attempting to locate it. Granted there is a lack of clarity around the notion of a motive, but on most any reasonable account a motive will be roughly an aim, goal or desire. And it is hard to see how one could attempt to locate an object without aiming to locate it, having the goal of locating it, or wanting to locate it.

So there are two morals we can draw from the analogy. One is that Global Search 0.1 does not follow from Local Search. The other is that Local Search is almost certainly true, whereas Global Search 0.1 is arguably false. These points matter because Velleman's theory of reasons rests on claims that are closely analogous to Local Search, whereas what he needs for his Kantian strategy is something analogous to Global Search 0.1. So if his theory of reasons delivers only the local kind of inescapability, then it does not deliver what the Kantian strategy needs. And that, recall, is what I aim to show in this paper.

To finish showing this I need to do two things. First I need to explain why Velleman's Kantian strategy requires the global kind of inescapability. And then I need to show that the theory of reasons delivers only the local kind.

Why then does the Kantian strategy need the global kind of inescapability? Why wouldn't local inescapability be enough?

Well, recall what the Kantian strategy is. The idea is to show that although every reason for action must rest on a current motive of the agent, there may be some reasons that rest only on a motive that is necessarily shared by all agents. For this to work, there must be a motive that is necessarily shared by all agents. In other words, the following claim must be true:

> **Global Action)** There is a motive such that, necessarily, anyone who does actions has that motive.

The reason local inescapability is not enough is that local inescapability does not entail global inescapability. We've already seen this for the case of searches. Now we need to see it for the case of action.

Just as one might search for one thing but not for another, so one might attempt to understand one's behavior in one respect but not in another. Perhaps I am currently making shrimp Alfredo, guided by a conception of making shrimp Alfredo, and I am also sighing, but not

guided by a conception of sighing. In Velleman's terms, my making shrimp Alfredo is an action, whereas my sighing is mere behavior. The difference, he says, is that I am making shrimp Alfredo in an attempt to behave in way I can understand, whereas although I am sighing, I am not sighing in an attempt to behave in a way I can understand.

I take it, though, that I am not making shrimp Alfredo in an attempt to make sense of all of my behavior. I am not, for instance, making shrimp Alfredo in an attempt to sigh in a way that makes sense. Indeed I may not even be aware that I am sighing. Rather, my aim in making shrimp Alfredo is to understand my behavior in respect of making shrimp Alfredo – my making it, or not making it, or making it this way or that, as the case may be. Or to put it another way, my aim is to exhibit shrimp Alfredo related behavior that I can understand. Thus if I stubbornly make shrimp Alfredo when it would make more sense make something else, or if I make it in the oven when it would make more sense to make it on the range, then my making it at all, or my making it in the oven, constitutes a failed attempt to exhibit shrimp Alfredo related behavior that makes sense. With this in mind we can formulate local inescapability for action as follows:

> **Local Action)** The motive of understanding one's φ-ing related behavior is inescapable for anyone whose φ-ing related behavior amounts to action.

And now we can see why, in the case of action, local inescapability does not entail global inescapability. It is because Local Action can be true even if there is no sphere of behavior that must amount to action for all agents. Suppose you have never heard of shrimp Alfredo, and are not currently engaged in any attempt to exhibit shrimp Alfredo related behavior that makes sense. That need not stop you being an agent, since you may be engaged in attempts to behave intelligibly in other respects. And for all Local Action says, it may be that every respect is optional in this way.

So Local Action does not entail that there is any motive that is necessarily shared by agents. For all it says, every motive may be contingent, in the sense that there could be an agent who lacked it. Global Action, on the other hand, says there is a motive that is not contingent in this sense. It says there is a motive that no agent could lack while still being an agent. Thus Local Action does not entail Global Action.

And for the same reason, Local Action does not entail the more specific version of Global Action that Velleman wants to emerge from the philosophical analysis of agency:

> **Global Action 0.1)** There is a motive such that, necessarily, anyone who does actions has that motive, and it is the motive of self-understanding

For suppose Global Action 0.1 did follow from Local Action. Then since Global Action 0.1 entails Global Action, Global Action would follow from Local Action. But we've just seen that Global Action does not follow from Local Action, so we have our contradiction.

Velleman's Kantian strategy consists in showing that there is there is a motive that no agent could lack while still be an agent, a motive whose presence is a prerequisite for being subject to reasons at all. And Velleman hopes to show this by showing that the motive of self-understanding is such a motive. His strategy, therefore, requires that both Global Action and Global Action 0.1 be true. But neither Global Action nor Global Action 0.1 follows from Local Action. So Local Action does not deliver what the Kantian strategy needs.

That was the first thing I needed to show. The second was that Velleman's theory of reasons, which I have here granted for the sake of argument, delivers only local inescapability. Given my reconstruction of Velleman's theory of reasons, this falls out fairly easily. I do grant, though,

that my reconstruction might be inaccurate, so that what I say here could be questioned on exegetical grounds.

Consider once more the two analyses on which the theory of reasons rests.

> **Analysis of Belief)** For S to believe that Q is for S to represent Q as true in an attempt to represent as true whatever really is true as regards whether Q

> **Analysis of Action)** For A's φ-ing to be an action is for A to φ in an attempt to understand her behavior as regards φ-ing

The former says the motive of truth about Q is inescapable for those who are in the business of forming beliefs about Q. It does not say there is any proposition that believers necessarily form beliefs about. The latter says the motive of understanding one's behavior as regards φ-ing is inescapable for those whose φ-ing is an action. It does not say there is any sphere of behavior with respect to which every agent's behavior must amount to action. There is nothing in Velleman's analyses of belief and action, then, that goes beyond local inescapability.

Nor do the theories of reasons that are built on these analyses go beyond local inescapability:

> **Theory of Reasons for Belief)** For P to be a reason for S to believe that Q (or to believe that not-Q) is for P to be an indicator of what would make for success in an attempt to represent as true whatever really is true as regards whether Q

> **Theory of Reasons for Action)** For P to be a reason for A to φ (or not to φ, or to φ this way or that) is for P to be an indicator of what behavior on A's part would make for success in an attempt to understand her behavior as regards φ-ing

The theory of reasons for belief analyzes reasons for belief on a topic in terms of what would make for success in an attempt at truth on that topic. It does not commit us to the idea that believers have any general motive of representing as true whatever is true.[13] It may be that humans sometimes do want there to be truths that they represent as true, much as humans sometimes want there to be objects that they locate. But nothing of the kind follows from Velleman's theory of reasons for belief.

Likewise, the theory of reasons for action analyzes reasons for φ-ing in terms of what would make for success in an attempt to understand one's φ-ing related behavior. It does not commit us to the idea that agents have any general motive of behaving in ways they will understand. Again, it may be common among humans to want there to be behaviors of theirs that make sense, and this could move them to pick a behavior that makes sense and do it, much as one might pick an object and set out to locate it. But nothing like this follows from the theory of reasons for action.

Velleman's theory of reasons, I conclude, delivers only local inescapability. What he needs for his Kantian strategy, however, is global inescapability. His theory of reasons, therefore, does not support his Kantian strategy.

That was what I set out to show. But it is worth emphasizing the limits of this conclusion. I have not questioned Velleman's account of action, or his theory of reasons for action. Nor have

---

[13] Velleman (2000: 252) makes a very similar point about belief. "What distinguishes believing a proposition from imagining it or supposing it," he writes, "is a more narrow and immediate aim – the aim of getting the truth-value of that particular proposition right … Belief is the attitude of accepting a proposition with the aim of thereby accepting *a* truth, but not necessarily with any designs on truths in general, or Truth in the abstract."

I argued that his Kantian strategy must fail. I have merely questioned his attempt to ground the Kantian strategy in the philosophical analysis of agency. The idea that there is some one motive that "provides our shared constitution as agents," as he puts it, simply cannot be derived from what he says about agency. For all I have said here, there may yet be such a motive. Perhaps its existence can be shown in some other way.

Now for a bit of diagnostic speculation. I've argued that Velleman can't spin the motive needed for the Kantian strategy from his philosophical analysis agency, but why would it seem that he could?

One possibility is that anyone who expects to engage in attempts of one sort or another is likely to want those attempts to be successful. Thus for example, if you figure that in the future there will be things you need to find, you will care whether your searches for those things succeed. And then if someone asks you, "Do you care whether you ever locate anything?" your answer must be yes, because if you never locate anything that means none of your searches will be successful, and who wants that? So perhaps there is a desire that all searchers will share, namely the desire not to fail in one's searches. And perhaps there is a desire that all believers share, namely the desire not to make failed attempts at the truth, and a desire that all agents share, namely the desire not to make failed attempts to behave intelligibly. So it appears that there is after all a way to derive a shared motive from the idea that beliefs and actions fall in the category of attempts.

But even if this is right, it does nothing for the Kantian strategy. The reason is that on the Kantian strategy, the motive that is necessarily shared by all agents must be one that can ground reasons for action. So the motive must be in a position to explain the actions for which it provides reasons. It must be possible, Velleman thinks, to come to do the action by a process of deliberation that starts from the motive. And that means the motive must be one from which particular attempts to behave intelligibly can be derived.

As we saw in our discussion of the motive of locating objects, there is a way of understanding that motive that allows it to play this role. If we take the motive of locating objects to be a general desire that there be objects that you locate, that desire can explain the motives of particular searches. For it can move you to pick some object and set out to locate it.

But the sort of desire we have just stumbled on is not like this. If you foresee that there will be things you need to find, and that is why you care whether you ever find anything, then the explanation runs from the anticipated searches to the desire, not the other way around. This is not a desire from which the motives of particular searches can be derived. You just have to let the searches arise as needed, and then execute them as best you can.

For the same reason, the analogous desires are not properly positioned to explain why one seeks the truth on some particular topic, or makes some particular attempt to behave intelligibly. If you foresee that there will be questions for which you seek true answers, and that is why you care whether you ever find the truth, that is not a desire from which the motives of particular beliefs can be derived. And likewise if you foresee that there will be things you do in an attempt to behave intelligibly, and that is why you care whether you ever behave intelligibly, that again is not a desire from which the motives of particular actions can be derived.

Perhaps, then, it is this phenomenon that explains the temptation to think the philosophical analysis of agency can support the Kantian strategy. We feel, reasonably enough, that all of us care about the success of our attempts. But then we mistakenly infer that the motives of particular attempts can be derived from this desire.

Another possibility is that we mistake a shared property for a motive. Local inescapability does entail that something is shared. For instance it does follow from Local Search that there is something all searchers share. It follows that they all have the property of being such that there is an object that they have the motive of locating. It does not follow, though, that there is any motive that they share. So perhaps we feel, rightly, that there is something all believers share, and something all agents share, but mistakenly infer that what they share is a motive.[14]

4   Conclusion

The Kantian strategy derives no support from the theory of reasons. But so what? How does this result affect Velleman's ambitions?

The answer may depend on which ambitions we have in mind. The one on which I have focused is the project of accounting for the objectivity of morality without positing a normative reality of which moral judgments can be true. The idea was to reconcile Williams's point that all reasons for action must be geared to present motives with an objectivist view of those reasons. The solution was to explain how the demands of practical reason could be "woven into the practical point of view." And the explanation was that if we look at what it is to act, we will find that there is a motive that must be present in anyone who qualifies as an agent. I've argued that even if we grant Velleman's analysis of action, and the accompanying analysis what it is to have a reason for action, we do not thereby find that there is any such motive.

But perhaps there are other ways of showing that the demands of practical reason are woven into the practical point of view. One possibility would be to stop looking for a single shared motive, and make do with Velleman's theory of reasons on its own. Could we then reconcile Williams's idea with the objectivity of practical reason?

It appears that we could not. For unless there is a necessarily shared motive, all motives are contingent for agents. And in that case we are forced to choose between the two options Velleman wanted to avoid. We must either bind reasons to contingent motives or sever the connection with current motives altogether.

Another possibility would be to stop looking for a motive that is constitutive of agency, and look instead for a motive that is reliably shared by human beings. Suppose Velleman were to accept that there could be agents who made only local attempts to behave intelligibly, and who lacked any more general aim from which the motives of those attempts could be derived. He could still claim that we humans are not such agents.

And in fact we know he would claim this, because he does. He distinguishes two senses in which the aim of self-understanding is inescapable. It is constitutively inescapable for us as agents, he says (137). But it is also naturally inescapable for us as human beings, in the sense that we humans can't help seeking to understand ourselves. Or rather we can only opt out of this aim temporarily, by losing ourselves in an activity or in daydreams, or by going to sleep (137).

What I am now suggesting is that Velleman could drop the idea that there is a motive that is constitutively inescapable for agents, and pin his hopes on the claim that there is a motive that is naturally inescapable for us as human beings. Perhaps he could argue that the demands of practical reason are woven, not into the practical point of view *per se*, but into the human point of view. That might be a way of combining Williams's internalism with a kind of objectivism about practical reason.

---

[14]I thank Gurpreet Rattan for this suggestion.

To make this out, he would need to show not only that humans can't help doing behaviors in attempts to behave intelligibly, much as they can't help forming beliefs, but that there is some one motive that humans can't help having. He would also need to show that this was a motive from which particular attempts to behave intelligibly could be derived. And finally he would need to show that there are reasons for action that rest on that motive and no other.

There is evidence that Velleman has such a program in mind. For he argues at some length that empirical work in social psychology supports the hypothesis that humans have a "drive toward self-understanding (17, 64).[15]"

Another possibility, though one to which Velleman is clearly less inclined, would be to rethink the commitment to Williams's conclusion. Does it really follow, from Williams's premises, that all reasons for action rest on current motives of the agent? One might rather reject this conclusion, instead of looking for a kind of objectivity that coheres with it.[16]

References

Clark, Philip (2001) "Velleman's Autonomism," *Ethics* 111: 580-93

Enoch, David (2011) "Shmagency Revisited," in M. Brady ed., *New Waves in Metaethics* (Basingstoke: Palgrave Macmillan) 208-33

Enoch, David (2006) "Agency, Shmagency: Why Normativity Won't Come from What is Constitutive of Agency," *Philosophical Review* 115, 169-98

Korsgaard, Christine (1986) "Skepticism about Practical Reason," *The Journal of Philosophy* 83, 5-25

Setiya, Kieran. (2003) "Explaining Action," *The Philosophical Review* 112:3, 339-393

Velleman, J. David. (2009) *How We Get Along*, Cambridge: Cambridge University Press

Velleman, J. David. (2006) *Self to Self: Selected Essays*. New York: Cambridge University Press

Velleman, J. David. (2000) *The Possibility of Practical Reason*, Oxford: Oxford University Press

Williams, Bernard (1995) *Making Sense of Humanity and Other Philosophical Papers 1982-1993*, Cambridge: Cambridge University Press

Williams, Bernard (1981) "Internal and External Reasons," in Williams, *Moral Luck*, Cambridge: Cambridge University Press, 101-13

---

[15]See also Velleman, "From Self Psychology to Moral Philosophy," in his (2006).

[16]Velleman (2000: 172 ff) considers Korsgaard's (1986) critique of Williams (1981). He reads her as severing the connection between reasons and current motives of the agent, and objects that she then faces a "burden of justification." For a reply to this objection, see my (2001).