

WHY COGNITIVISM?

[Preprint; final version forthcoming in the *Canadian Journal of Philosophy*]

ABSTRACT Intention Cognitivism – the doctrine that intending to V entails, or even consists in, believing that one will V – is an important position with potentially wide-ranging implications, such as a revisionary understanding of practical reason, and a vindicating explanation of 'Practical Knowledge'. In this paper, I critically examine the standard arguments adduced in support of IC, including arguments from the parity of expression of intention and belief; from the ability to plan around one's intention; and from the explanation provided by the thesis for our knowledge of our intentional acts. I conclude that none of these arguments are compelling, and therefore that no good reason has been given to accept IC.

I

'Intention Cognitivism' is the view that intending to V entails, or even consists in, belief. Roughly, a future-directed intention to V is or entails believing that one will V ; a present-directed intention to V is or entails believing that one is V -ing. This paper urges the rejection of Intention Cognitivism, as it is a thesis lacking adequate support. The main considerations usually marshaled in support of the thesis are explored, and the conclusion is pressed that they are unconvincing.

The idea that intention *entails* belief will henceforth be labeled 'Weak Intention Cognitivism', to be distinguished from 'Strong Intention Cognitivism', which will designate the thesis that intention *consists in* belief. Versions of Intention Cognitivism weak and strong have been put to use by various philosophers along the years with wide-ranging implications. First is the intrinsic significance of tracing any structural relations that hold between intention and belief, when seeking a theoretically adequate picture of mental economy. But the acceptance or otherwise of Intention Cognitivism also impinges on positions within a host of other debates in moral psychology and normative theory. For example, Gilbert Harman suggests that Weak Intention Cognitivism (WIC) offers the characteristic mark of intending, in contradistinction to other practical attitudes such as desires, hopes, wishes, and so on (1976: 432). If correct, Harman's view would establish how intention differs from other practical attitudes, and indeed *that* it differs from them; it

would call into question the familiar view on which intending to V is reducible to a desire to F coupled with some instrumental belief about how V -ing promotes F .¹

Harman takes WIC further to anchor a unified picture of reasoning across the practical and theoretical domains. Theoretical reasoning for Harman is an attempt to increase overall explanatory coherence among one's beliefs. And since intention entails belief, the same also holds for practical reasoning: WIC 'brings these two sorts of reasoning under the same principle' of belief consistency (1976: 435). Harman thus deploys WIC in the service of a unified understanding both of the nature of reasoning, and of the *principles* that govern it. Others have argued in a broadly similar spirit about the principles or requirements of *rationality*. They suggest that at least some requirements of rational coherence involving intentions actually derive from corresponding requirements of belief consistency.² Take for instance the rational requirement not to intend to V while intending to F and believing that one cannot both V and F . One who has this practically incoherent combination of attitudes is said to have a corresponding combination of inconsistent *beliefs* – viz. a belief that one will V , a belief that one will F , and a belief that one will not both V and F .³

Cognitivists offer also a parallel explanation of instrumental rationality. Kieran Setiya proposes the following cognitivist version of the requirement for means-ends coherence:

[ME-Cognitive] Rationality requires that [if you intend to V , and you believe that you will not V unless you M because you now intend to M , then you believe that you will M because you now intend to M].⁴

Setiya claims that ME-Cognitive can be used to derive the corresponding practical requirement:

¹ For versions of the familiar view, see Davidson (1963), and more recently Ridge (1998), and Sinhababu (2012). That intention may be reducible to desire combined with belief that one will V will remain a live possibility.

² See for example Wallace (2001), Ross (2009), and Setiya (2007).

³ Cognitivist explanations of rational requirements are typically founded on *Strong* Intention Cognitivism, although arguably the weak version would be sufficient for the arguments to go through. The point matters little here, however, since the present target is the weak version which is entailed by the strong. If the arguments that follow succeed, therefore, neither version is available to prop up the cognitivist understanding of practical rationality.

⁴ Setiya (2007b: 663-673).

[ME-Practical] Rationality requires that [if you intend to V , and you believe that you will not V unless you M because you now intend to M , then you now intend to M].

Given WIC, the first clause of ME-Cognitive can be replaced with the belief that you will V , yielding a purely theoretical requirement: one who violates ME-Cognitive would be violating epistemic closure. If successful, the above offers a revisionary understanding of practical requirements of rationality. It may also indicate how to resolve the on-going controversy over whether, and if so why, those requirements are *normative* or reason-giving:⁵ if requirements of practical coherence are explained by requirements of belief consistency, then presumably the normativity of the former is likewise explained by the normativity of the latter.

These are some of the implications of Intention Cognitivism for moral psychology and normative theory. One other possible implication is a vindicating explanation of the phenomenon of ‘Practical Knowledge’. Inspired by Anscombe’s (1957) controversial remarks, some try to account for the peculiar knowledge we supposedly have of our intentional acts on the basis of Intention Cognitivism. Anscombe argues that one has knowledge of what one is doing when acting intentionally (and perhaps also of what one will do when intending to act) that is peculiar in being immediate and not derived from observation, inference, or any other kind of evidence in support of the belief that one is V -ing. And some cognitivists seem to think that Intention Cognitivism explains how such otherwise mysterious knowledge can in fact be achieved: intention entails or consists in belief, and under favourable circumstances the belief amounts to knowledge.⁶

The above lists some of the theoretical payoffs afforded by Intention Cognitivism. In what follows, however, it will be argued that these payoffs are unavailable: the grounds for accepting WIC, and a fortiori also SIC, turn out to be specious. That conclusion will be reached by unfavorably examining the main arguments commonly offered in support of the thesis. Of course, payoffs may themselves be adduced as support: if Intention Cognitivism offers an attractive explanation of some extant phenomenon, then by a

⁵ See, for example, Kolodny (2005), Broome (2005), and Raz (2005).

⁶ Velleman (1989) & (2007), and Setiya (2008) & (2007b). This brief statement of Velleman’s and Setiya’s views is elaborated in sec. III.

kind of inference to the best explanation, the thesis itself should be believed. In fact, something like this pattern of inference from practical knowledge to Intention Cognitivism seems to be invoked by those who endorse both theses, and is handled accordingly below: it will be argued that Intention Cognitivism cannot actually explain knowledge of intentional action (Section III).⁷ Before then, the following section (II) aims to cast doubt on Intention Cognitivism by undermining two other central motivations for believing the thesis: one based on how intentions and beliefs can be expressed using the same sentences; and the other based on how intentions allow for planning. The paper's conclusion is thus that no adequate grounds for accepting WIC are currently available. If cogent, the appropriate response to this conclusion seems to be withholding judgment on the view, unless and until more compelling grounds emerge

II

One central source of motivation regularly invoked by Intention Cognitivists comes from the parity of verbal expression of intention and belief. 'I am going to V ', 'I shall V ', and 'I will V ' are all instances of canonical expressions of prospective intention using indicative sentences, used also to express beliefs; similarly for 'I am V -ing' as an expression of intention-in-action. And the thought seems to be that WIC best explains the parity of expression: it is *because* intention entails belief that one may express both attitudes with one and the same utterance. Of course, expression by means of the same *sentence* does little by itself to establish that both attitudes are actually expressed: at least the argument must show that the sentence in question is also uttered with the same force. The point is therefore often put by saying that intentions are canonically expressed by the *assertoric* utterance of 'I am going to V ' and the like. This supposedly ensures that in addition to intention, the speaker also betrays a corresponding belief.

Often attributed to Anscombe (1957), the argument from parity of expression is heralded by Kieran Setiya as 'the most powerful argument for [WIC]'.⁸ Since this argument takes the form of an inference to

⁷ The other payoffs of WIC mentioned above, to do with how it may account for the nature of, and the principles governing practical reason may also be deployed as part of an IBE in support of WIC. These further possibilities are not discussed here. For arguments against the cognitivist explanation of rational requirements of coherence, see Bratman (2009), Brunero (2009), and Broome (2013: 163-6).

⁸ Setiya (2007b: 34). Others who endorse the argument in some form or another include also Velleman (2007: 206-7) and Grice (1971), among others. Grice sets out to show that 'to say of someone that he intends to do A (without qualification) is ... to imply or suggest that he does not think it doubtful whether he will in fact do A' (1971: 264).

the best explanation, doubts over its cogency may have at least two sources. First, one may deny that there is in fact a genuine phenomenon here for WIC to explain; this is the first point raised immediately below (II.1). The assumption that intentions are expressed by asserting is called into question, thereby questioning also the purported parity with expressions of belief. The second doubt raised in this section (II.2) allows that parity of expression may be real, but puts forward a competing explanation for this fact which discredits the suggestion that WIC *best* explains the phenomenon. It is argued that even if intentions *are* generally expressed by asserting, the presence of a corresponding belief may be conversationally implicated rather than entailed.

II.1 Parity of expression?

To start with, note that the linguistic evidence for a robust correlation between expressions of intention and belief is far from decisive. Sentences such as ‘I am going to V ’ and the like are no doubt a natural and ubiquitous way of expressing one’s intention; but they are hardly the *only* way. Often enough, one expresses one’s (future) intention by saying e.g. ‘I intend to V ’. And the latter of course is not used to express any belief that one will V . It might therefore be claimed that at least those intentions expressed by ‘I intend to V ’ do not exemplify the parity of expression WIC purports to explain. (Indeed, as discussed later, expressing one’s intention by uttering ‘I intend to V ’ seems to be a means of signaling that one is not confident, i.e. lacks the belief, that one will in fact V).

Further, even ‘I am going to V ’ and the like are not used *exclusively* to express intention but sometimes also other attitudes, such as hope (e.g. ‘I’m going to pass the exam’, uttered by the unprepared student attempting to reassure himself), wish or desire (‘I am going to be prime minister’, uttered by the five-year-old), incredulity, or even (uttered in a quivering or timid tone of voice) fear. Hence, there need

To support this claim, Grice offers the following ‘imaginary *conversation*’ [emphasis added]:

- X. I intend to go to that concert on Tuesday.
- Y. You will enjoy that.
- X. I may not be there.
- Y. I am afraid I don’t understand.
- X. The police are going to ask me some awkward questions on Tuesday afternoon, and I may be in prison by Tuesday evening.
- Y. Then you should have said to begin with, ‘I intend to go to the concert if I am not in prison’, or ... ‘I should probably be going’, or ‘I aim to go’, or ‘I intend to go if I can’.

be nothing uniquely common to expressions of intention and belief that calls out for explanation. And in fact, when a speaker expresses any of the other attitudes above by uttering ‘I am going to V ’, we are not tempted to suppose that she must also believe she will V ; nor, conversely, are we tempted to suppose that a desire, wish, hope etc. are present whenever a speaker expresses her *intention* by uttering ‘I am going to V ’ and the like. Why, then, should we take the canonical expression of intention in particular to indicate the presence of a corresponding belief?⁹

The cognitivist may insist in response that there is a more robust correlation between expressions of intention and belief than between any of the other attitudes mentioned above. For ‘I am going to V ’ and the like represent the *standard* or canonical way of expressing both an intention to V and a belief that one will V , unlike the other attitudes towards V -ing which are more rarely expressed by these sentences. This is what uniquely singles out the correlation between expressions of intention and belief as a striking fact that calls out for explanation – an explanation that WIC is best placed to provide.

Still, this response falls short of establishing parity of expression, as will now be argued. Go back to the innocuous, incontrovertible observation from which the target argument sets out: ‘I am going to V ’ (‘I am V -ing’) and its cognates are typically used to express either an intention to V or a belief that one is going to V (is V -ing). When used to express the former, are they uttered with assertoric force which would imply the existence of the latter? Proponents of WIC typically assume an affirmative answer without argument. The assumption may be unproblematic if no alternative proposal could be plausibly maintained. But that does not seem to be the case. Versions of *expressivism* have famously been deployed to produce a viable revisionary understanding of utterances bearing surface propositional form, in such various and different discourses as ethics, aesthetics, humour, knowledge, causation, probability, and even logic.¹⁰ One such type that is particularly relevant here is expressivism about mentalistic attributions or avowals. ‘I’d like an espresso’, ‘I’m afraid of spiders’, ‘I have a splitting headache’, and so on, are all sentences which, if sincerely uttered, seem to *report* mental attitudes (desires, sensations, feelings, etc.) self-

⁹ The discussion in the text of the canonical expression of prospective intention may not carry over to the expression of intentions with which one is acting.

¹⁰ For a useful survey of, and references to, the growing literature on the different forms of expressivism, see Bar-On and Sias (2013).

ascribed by the speaker. But a familiar idea is that avowals actually serve to directly *express* rather than report the mental condition the speaker is in.

Avowal-expressivism is motivated on grounds independent of the present discussion, viz. primarily the prospect of vindicating the peculiar epistemic standing of avowals. The unparalleled epistemic security enjoyed by avowals is a very widely recognized phenomenon. Avowals are typically *immediate*, i.e. not made on any evidential base (observation, inference, etc.). Further, avowals are typically exempt from ordinary forms of epistemic assessment: they are presumptively regarded as true, and only rarely contested, criticized, or rejected. Indeed, avowals are epistemically privileged *even* in comparison with other first-personal non-mental ascriptions ('My legs are crossed'; 'There is a cube in front of me'; 'I'm a patient person'), which either *are* typically based on evidence, or more readily challengeable (or both).

A chief advantage of avowal-expressivism is that it accommodates the epistemically privileged position of avowals without recourse to what many regard as the problematic suggestion of a unique epistemic base or method, such as introspection. Briefly, the central idea is that avowals are on a par with natural expressions, e.g. groans and winces, in speaking directly from or giving voice to one's mental condition, rather than reporting it. And this is what explains their distinctive and privileged epistemic standing: similarly to natural expressions, avowals are ungrounded in any evidential base and are typically immune to epistemic assessment.¹¹

Expressions of intention are arguably continuous with, if not a subset of, avowals. 'I am going to \mathcal{V} ', 'I will \mathcal{V} ', and 'I am \mathcal{V} -ing' may not obviously seem to be making self-ascribed reports of *intention* (unlike 'I intend to \mathcal{V} '); but they clearly seem to be reporting some such mental item – if not intention, perhaps a plan or a commitment or indeed an action, either ongoing or future, which is at least partly mental. Hence the present case affords a natural candidate for the attitude being expressed (unlike expressivist treatments

¹¹ For discussion, see for example Wright (1998), Bar-On and Long (2001), and Bar-On (2004). Bar-On's sophisticated Neo-Expressivism incorporates a distinction between 'expression' in the sense of the *act* or process of expressing some sentence ('expression_A'), and 'expression' in the sense of the *product* of that process ('expression_P'). Considered as act, Bar-On insists that avowals express_A by directly giving voice to one's mental condition; but considered as product, avowals are indicative sentences expressing_P a proposition. The distinction allows Bar-On to uphold the epistemic security of avowals while circumventing some of the traditional obstacles facing avowal-expressivism, to do with how it might account for the salient syntactic and logico-semantic continuities between sentences in the expressivist domain and sentences in other, non-expressivist domains. On Bar-On's account, any avowal made by a declarative sentence expresses_P a proposition so that, once again, there need be no striking correlation between expressing intention and belief.

of some other types of claim, e.g. causal ['the glass broke because it was struck'] or conditional ['if the glass is struck, it will break'], where it is *prima facie* harder to see what the relevant attitude might be). Further, expressions of intention enjoy the sort of privileged epistemic status typical of avowals described above. Treating avowals of intention along expressivistic lines is therefore no less plausible than treating other avowals in this way. And, as already noted, such treatment is motivated independently of the present discussion of WIC, by how it upholds the epistemic security of avowals.

The above remarks do not pretend to have established anything resembling a workable expressivistic conception of intention-avowals; they barely make a start in that regard. *A fortiori*, they can hardly claim to have refuted the assertoric conception. But that is anyway not their purpose. More modestly, they are meant to clear ground for the *possibility* of expressivism about intention-avowals. This is enough to invalidate the presumption that intentions are expressed by means of asserting propositions. And with that presumption invalidated, little is left of the parity of expression the cognitivist originally seizes on. The mere fact that intention is expressed – *even standardly* – by declarative sentences hardly indicates the existence of a striking correlation with expressions of belief: declarative sentences are (standardly) used to express any number of attitudes without also (standardly) expressing beliefs – cf. 'I'd like to know what time it is' used to ask a question, or 'So happy to see you!' spontaneously expressing joy, to cite just a few examples. This is the first doubt raised here about the support WIC supposedly gains from how it explains the verbal expression of intention. It is unclear that parity of expression is in fact a genuine phenomenon in need of explanation.¹²

Of course, as noted, the above does not rule out the possibility that intentions *are* typically expressed by assertions, and hence that there is in fact a striking correlation with expressions of beliefs which WIC is uniquely positioned to explain. Nevertheless, the *prima facie* plausibility of the alternative expressivistic view shows that as it stands – i.e., prior to demonstrating that expressivism should be rejected – the argument from parity of expression is at best inconclusive.¹³

¹² It might be thought that the speech act of expressing intentions serves a function, e.g. to facilitate *planning*, which requires that one have the corresponding belief. This suggestion is addressed below.

¹³ I am grateful to an anonymous referee for urging me to clarify the dialectical target of my argument.

II.2 Entailment or conversational implicature?

But let all that pass. It will now be argued that even if intentions *are* expressed by asserting, this does not underwrite compelling support for WIC. For there is available a plausible alternative explanation of the parity with expressions of belief, which renders doubtful the suggestion that WIC offers the best explanation of this phenomenon. As already noted, proponents of WIC often cite Anscombe (1957) when relying on the argument from parity of expression. Careful attention to one of her examples reveals why the assertoric expression of intention may *seem* to support WIC – and why it actually does not (no attempt will be made to determine Anscombe’s own view on the matter):

Nor can we say: But in an expression of intention, one isn’t saying anything is going to happen! Otherwise, when I had said: ‘I’m just going to get up’, it would be unreasonable later to ask: ‘Why didn’t you get up?’ I could reply: ‘I wasn’t talking about a future happening, so why do you mention such irrelevancies?’¹⁴

Saying ‘I’m just going to get up’ generates an expectation in the hearer that the speaker will in fact get up; that is why the reply Anscombe imagines is absurd. The expectation is plausibly based on a further expectation, to the effect that the speaker herself expects to get up, i.e. that she believes she will get up. Now does this entail that the speaker in fact believes she will get up? Not if her assertion can occasionally be defective in some way (not necessarily by being *insincere*). An infelicitous assertion will generate the very same expectation adduced as evidence for WIC, while failing to imply that the expectation is confirmed. In short, that one believes one will get up may be *merely* conversationally implicated.¹⁵

The cognitivist may insist in reply that if the speech act is defective in violating the belief rule of assertion, the speaker does not in fact express an intention to *V*. In her mouth, ‘I am going to *V*’ or ‘I am *V*-ing’ actually express some mental condition that falls short of intention – a hope, wish, attempt, or whatever. At this point the threat of deadlock looms. To overcome it, a procedure will now be outlined

¹⁴ Anscombe (1957: 92).

¹⁵ Something like this point is noted in passing by Davidson (1980: 91). The ensuing discussion in the text might be read as an elaboration and defense of Davidson’s brief comment. A different possibility, not explored here, is that one’s believing that one will *V* is a *presupposition* of the sentence expressing one’s intention to *V*.

which arguably confirms the above hypothesis that beliefs are merely conversationally implicated by expressions of intention.¹⁶

The procedure takes a cue from Moore's Paradox, which famously involves the absurd utterance '*p*, but I don't believe that *p*'. Moore's proposition may be perfectly true, and still there is something deeply odd about asserting it; hence the air of paradox. The oddity comes at least in part from expressing belief in some proposition, and then going on to deny believing that same proposition. The proposal here will be that if intention entails belief as the cognitivist maintains, a parallel to Moore's original paradox should arise. A proposition which expresses an intention to *V* and goes on to deny believing that one will *V* (for prospective intention) or that one is *V*-ing (for intention-in-action) should come out as paradoxical. The second clause of '*p*, but I don't believe that *p*' reports the absence of an attitude that the first clause implies is present. Similarly, the second clause of 'I am going to *V* but I don't believe I will *V*' likewise reports the absence of an attitude that the first clause implies is present.¹⁷

The modified Moorean proposition 'I am going to *V* but I don't believe I will *V*' is not fully analogous to the original. For one thing, different propositions occur in each clause ('I am going to *V*'; 'I will *V*'). But this different structure was chosen merely to make salient that the proposition is a modification, not an instance, of Moore's original paradox. As long as the reader can hear the first clause of 'I am going to *V* but I don't believe I am going to *V*' as expressing intention and not belief, she is free to substitute it without damage to the procedure it will serve. A more serious difficulty arises when intention-in-action is concerned. It is very hard not to hear 'I'm building a house but I don't believe I'm building a house' as simply an instance of Moore's Paradox. And it is not obvious that there is a suitable alternative formulation.¹⁸ The best candidate seems to be 'I'm *V*-ing but I don't believe I'm managing to *V*'. If this pulls too far from Moore's original for a clear verdict on present intention to emerge, still the

¹⁶A standard test for whether a sentence is conversationally implicated considers whether the implication is cancelable by the speaker. Cancelability is discussed later in this section.

¹⁷The literature sometimes mentions, alongside Moore's original 'ommissive' proposition, also its 'commisive' counterpart '*p*, and I believe that not-*p*'. To save words, only the ommissive version of the modified Moorean proposition will be discussed here.

¹⁸I intend to be building a house but I don't believe I'm building a house' sounds cumbersome and forced in virtue of its first clause. Moreover, the class of intentions-in-action as such cannot plausibly be understood as intentions *to be V-ing*. For many, perhaps most, intentions do not match this schema. These are intentions to actually complete some task or achieve some end (e.g. an intention to poison the archduke, or to clean the house), rather than to be *en route* to completing or achieving it. Finally, 'I intend to *V* but I don't believe I'm *V-ing*' seems unhelpful since it fails to mark out the first clause as applying in particular to present, rather than future, intention.

point about *prospective* intention stands: the following procedure demonstrates that parity of expression does not support WIC, at least as a thesis about the (pervasive) class of prospective intentions. (In fact, as explained in section III below, cognitivists tend to regard WIC as the feature that *unifies* prospective and present intention. This line of thought is undercut if WIC turns out to be unmotivated with respect to the former type). In any event, the implications of the following discussion will be pointed out for future and present intentions alike.

A further disanalogy between Moore's original paradox and its modified version is that the latter refers to *two* different attitudes and not just one. But that should not impugn its absurdity. To verify, consider 'Ah, if only Arsenal win the league! But I don't want Arsenal to win the league'. Hoping entails desiring, and hence this proposition is a Moorean absurdity. Its first clause expresses hope that some state of affairs will obtain, while the second denies wanting the same state of affairs to obtain.

Our primary question is whether expressing an intention by asserting 'I am going to V ' and the like implies or merely conversationally implicates that one believes that one will V . Under the present procedure, the question becomes whether the absurdity of the Moorean proposition can be fully explained at the level of speech act, or whether the explanation runs deeper to the level of mental economy. It is a standard constraint on any explanation of Moore's Paradox that it encompass both levels. Any adequate account of what exactly makes ' p , but I don't believe that p ' paradoxical should recognize the oddness both of asserting *and* of judging that proposition.¹⁹ Similarly, if 'I am going to V but I don't believe I will V ' is indeed a Moorean absurdity, this proposition should come out paradoxical both when considered as the object of assertion, and when considered as the object of a pair of mental items: a *decision* (to V), and the absence of belief (that one will V). If this constraint is not met, the present procedure suggests that the proposition is not in fact a genuine Moorean absurdity, contrary to what the verbal argument for WIC predicts. Consequently, the argument fails.

¹⁹ See for example Heal (1994), who lists as a 'first condition on a solution' to the nature of Moore's Paradox 'that it must be of adequate generality to explain the oddness of both thought and assertion' (1994: 6). Various different accounts of Moore's Paradox have been proposed, broadly dividing into suggestions that judging or asserting the proposition would indicate a severe form of *theoretical* or alternatively of *practical* irrationality. For a useful recent survey, see the introduction to Green and Williams (2007). Green and Williams also count the oddness of both assertion and judgment as a constraint on any account of Moore's Paradox (p. 10).

Is, then, our modified Moorean proposition clearly absurd when considered as the object of a decision coupled with the absence of a corresponding belief? Rehearsing briefly a host of familiar cases suggests that, unlike with Moore's original paradox, the answer is *unclear*. For example, Anscombe describes a man who is "as certain as possible that he will break down under torture, and yet determined not to break down" (1957: 93). Similarly, when the Polish cavalymen galloped towards the German tanks in 1939, they plausibly intended to stop the Germans from advancing; but it is unlikely that they believed they will succeed. Cognitivists might reply that the protagonists of these cases are not in fact intending to \checkmark . Rather, they are more naturally thought of as having the qualified intention to *try* to \checkmark . If the response is adequate, the above cases are not actually instances of non-absurd decisions to \checkmark coupled with lack of belief that one will \checkmark . The reply could be bolstered by suggesting that we hesitate to attribute intention in these cases precisely *because* a corresponding belief is absent.

However, the response faces trouble handling some other cases that do not plausibly involve intentions to try. Thus Bratman (1987: 38-9) famously describes a person who intends to stop at the bookstore on her way home. But being prone to forgetfulness, she does not believe that she will indeed stop; once she gets on her bicycle and starts pedaling, she knows she is more likely to head straight home without stopping. As described, the case does not seem to be one in which the agent intends to try to stop at the bookstore. As Richard Holton says, "it is not as though the [bookstore] is currently under siege, and you are sceptical about your abilities to get through the defences" (2008: 29-30). Moreover, even if this *were* a case of one's intending to try to stop, one cannot be described as having the corresponding belief that one will try to stop. For any doubts as to whether one will indeed stop will carry over to the attempt to stop.^{20, 21}

Much the same point about intention-*in-action* is illustrated by Davidson's well-known carbon copier (1980: 92) – the person pressing hard against a page with the intention of making ten carbon copies, even though he does not believe he is making ten copies as he doubts the impression is actually going through. Once again, as with prospective intention, the reply may be given that the man is merely pressing with the

²⁰ As will become clear, the point of rehearsing these familiar cases is to suggest that their correct interpretation is hard to make out, rather than to reinforce the interpretations advocated by their authors.

²¹ This is not to suggest that one's doubts over actually stopping would also carry over to doubts over trying to *remember* to stop. But trying to stop involves more than just trying to remember to do so. And according to IC, if one's intends to try to stop, one's intention should entail a corresponding belief that one will actually try to stop.

intention to *try* to make the copies. However, again, trying seems out of place here, too. For the case may be such that the man knows from experience that he *can* make the copies if only he presses hard enough; he is simply unsure whether he is doing so.

Any connection that *does* hold between one's intention and one's cognitive states may be captured by weaker conditions than WIC. Even if intending to V entails neither believing that one will V nor believing that one is V -ing, still it may entail some related doxastic constraints such as believing that one can V , or believing that one might V , or not believing that one cannot V , or not believing that one will not V . Any of the above, or some combination thereof, may be well-placed to account for the relation between intention in belief in a way that is compatible with the above cases as described.

To repeat, the point of rehearsing these familiar cases is *not* to establish their force as decisive counterexamples to WIC. The point is merely to highlight the *possibility* of interpreting them in a way that runs counter to WIC. In Bratman's case and its variants, doubts over the stability of one's intention translate to doubts over its successful execution due to forgetfulness, lack of self-belief etc., while the intention remains in place. That is a possible diagnosis.²² The diagnosis falls short of conclusive verdict, and the cognitivist will urge that we resist it. But for present purposes, a *prima facie* plausible interpretation is all that is required. For it casts doubt over the thought that explanations of the oddness of 'I am going to V , but I don't believe I will V ', or 'I'm V -ing but I don't believe I'm managing to V ', *must* extend to the level of mental economy, as with genuine Moorean absurdities. A plausible explanation could instead be given which is limited to the occurrence of false implicatures when such sentences are uttered, thus countervailing any convincing support WIC claims from the assertoric expression of intention.

The conclusion is bolstered by considering the standard test of *cancelability*. Grice famously observes that if a sentence is merely conversationally implicated and not entailed by another sentence, the implication can be canceled by the speaker.²³ Hence an utterance which explicitly removes the hearer's expectation that one believes that one will V as one intends should not sound inappropriate. Consider then the sentence: 'I am going to scale this mountain but I don't believe I'll make it – it's going to be very

²² Holton (2008: 31) reaches a similar conclusion about Bratman's case.

²³ Grice (1989), ch. 2: 'Logic and Conversation', and ch. 17: 'Presupposition and Conversational Implicature'.

hard'. Voicing doubts that one is able to scale the mountain should successfully remove the expectation that one believes that one will scale the mountain. Does it?

If the cancellation seems to fail as the above utterance sounds inappropriate, this may be explained by a number of pragmatic factors. First, 'I am going to scale this mountain but I don't believe I'll make it' is easily heard as a straightforward *instance* of Moore's Paradox, which is not how it should be interpreted here. Second, and relatedly, the utterance may seem to violate something like Grice's Maxim of Quantity: one is in a position to utter e.g. the more informative 'I *intend* to scale this mountain, but ...', which would circumvent any confusion over whether one is expressing belief or intention by uttering 'I am going to scale this mountain, but ...'.²⁴ Finally, a further reason why the original utterance may seem awkward is that 'I don't believe I'll make it' is easily heard as equivalent to 'I believe I won't make it', which avows or reports the belief that one will not scale the mountain. But the argument developed here against Intention Cognitivism does *not* assume that intending to V may entail believing that one will not V , and hence the two possibilities should be kept firmly apart. Bearing all these complicating factors in mind, a more perspicuous utterance to consider may be: 'I intend to scale this mountain, but I suspect I might not make it – it's going to be very hard'. And this utterance does not seem inappropriate.

Having seen that an intention-expressing-assertion of 'I am going to V ' and the like may only conversationally implicate that one believes that one will V , we can use this observation to cast doubt over another consideration sometimes taken to support Intention Cognitivism. David Velleman asks: 'If I am agnostic as to whether I will be in Chicago on Tuesday, why should anyone plan or act on the assumption that I will be there?' (2007: 206).

The reply suggested here is by now no doubt obvious. Others will know that Velleman intends to be in Chicago on Tuesday because he *tells* them. And it is plausibly his telling, not his intending, which gives them grounds to plan or act on the assumption that he will indeed be there. The possibility of others planning around one's intention does not require that intention entail belief.²⁵

²⁴ Saying 'I intend to V ' rather than 'I am going to V ' may sometimes serve to weaken one's commitment to V -ing, which is itself a complicating factor.

²⁵If, as in the scenario Velleman imagines, the implicature that Velleman believes he will make it to Chicago is actually false, then his audience will be making plans on a misleading basis. According to the interpretation suggested in the text, this is because Velleman has failed to cancel the implicature. Cancellability is discussed earlier in this section.

The further question might arise, why should *oneself* plan or act on the assumption that one will be in Chicago on Tuesday if one is agnostic. That much cannot be explained by appealing to rules of assertion or implicatures; one does not communicate one's intention to oneself. At this point it is worth noting that Velleman's talk of *agnosticism* may be oversimplified.²⁶ Even if one does not believe that one will be in Chicago (not: believes that one will not be in Chicago), still one could believe that one *might* be there, or that there is a reasonable chance that one will be there, and so on. And having such beliefs should be enough, at least to make tentative plans.²⁷

III

A final source of support for Intention Cognitivism to be examined here comes from how it supposedly underwrites a vindicating explanation of the knowledge one has of one's intentional acts. If this so-called 'practical knowledge' is a genuine phenomenon, and moreover one that is best explained by Intention Cognitivism, the latter could claim support by a kind of inference to the best explanation. Something like this line of thought is explicit in some cognitivists, for example Setiya (2012).²⁸ Others, e.g. Velleman (1989, 2007), likewise rely crucially on Intention Cognitivism when explaining the possibility of practical knowledge.²⁹

The knowledge in question is thought of as a kind of self-knowledge, similar to knowledge of one's own mental attitudes (including one's intentions) in its unique and privileged epistemic standing (see pp. 6-7 above for more details). Here is Anscombe illustrating this point:

Say I go over to the window and open it. Someone who hears me moving calls out: What are you doing making that noise? I reply: 'opening the window' ... But I don't say the words like this: 'Let me

²⁶ Cf. Holton (2008: 31). In that paper, Holton develops an interesting notion of partial intention that involves tentative plans backed up by other contingency plans.

²⁷ The *locus classicus* of the planning theory of intention is Bratman (1987). Interestingly, Bratman himself denies that intention entails belief, as his discussion of the bookstore and other examples makes clear. Velleman's discussion in his (2007) sets out from an objection to Bratman's view.

²⁸ See in particular Setiya (2012: 193-197).

²⁹ Both Velleman and Setiya argue that intention consists in a kind of belief. Setiya's view is actually more complex than both WIC and SIC. He argues that intending to V entails either a belief that one will V , or that one has is more confident that one will V than one would otherwise be, or that one is taking some means towards V -ing (Setiya 2008: 391). It is not altogether clear whether the latter two beliefs are sufficient to support Setiya's claim that intending to V consists in knowing that one *will* V . For discussion, see Paul (2009a) and Setiya (2009). And see more below.

see, what is this body bringing about? Ah yes! The opening of the window'. Or even like this 'Let me see, what are my movements bringing about? The opening of the window'. (1957: 51)

But there is something deeply puzzling about the intuitive appearance of such knowledge as stated by Anscombe, which makes it a highly contentious idea. Unlike believing, desiring, and so on, to know that one is intentionally acting is to know that something is taking place in the environment *outside* one's mind. It thus seems to outstrip mere self-knowing, and must after all rely on ordinary methods of gathering evidential support – observational, inferential, or whatever. One needs some assurance that one's bodily movements are in fact successfully carrying out the action one intends to preform.³⁰

Intention Cognitivism is thought of as a crucial step towards upholding the intuitive picture of knowledge of action by first ensuring the presence of *belief*, which amounts to knowledge under favourable conditions. When one intends to V , one believes that one is going to V . When the intention is being realized, the belief becomes contemporaneous: one believes that one is V -ing. Thus Velleman writes:³¹

When one says 'I am going to take a walk', one lets the hearer know that one is going to take a walk. One's assertion is meant to provide the justification in virtue of which the hearer then knows that one is going to take a walk, and it is meant to provide that justification by virtue of expressing one's own knowledge to the same effect. Hence, an expression of intention must at the same time be an expression of knowledge – of something known, in other words, by being the content of intention.

But why suppose one's belief that one is going to take a walk amounts to knowledge, if it is not supported by evidence? Velleman (1989: 56-57) argues that the belief-in-intention need not be formed on the basis of evidence since once the belief *is* formed, it is guaranteed to enjoy post-hoc justification: once one forms the intention to V , one is justified in believing that one will V because one is aware of the intention and is inclined to act as one intends. Hence on Velleman's view, one is entitled to 'jump to the conclusion' that one is going to V since the conclusion is self-supporting. And similarly for knowing that one is currently V -ing: one's belief that one is now V -ing need not antecedently be based on any evidence, since once the

³⁰ For this point, see e.g. Donnellan (1963).

³¹ Velleman (2007: 201).

corresponding intention to be *V*-ing now is in place, the belief is guaranteed to be justified. Setiya questions Velleman's account, claiming that one's belief is justified not only post-hoc, but also in advance of forming the belief-cum-intention. On Setiya's view, the justification is supplied, though not derived from, one's *knowledge how* to perform the intended action (Setiya 2008: 401-6).

Most critical discussions of 'practical knowledge' question the claim the one could know that one is *V*-ing as one believes, absent any evidential or inferential base from which it is derived.³² The focus in what follows will be different. As Velleman and Setiya make clear, the cognitivist strategy for vindicating self-knowledge of contemporaneous action entails self-knowledge of *prospective* action. But the latter is an implausible upshot that is untrue to the phenomenon, rendering doubtful any claim that the cognitivist strategy offers the best explanation thereof. Whatever truth may be in the idea of practical knowledge, it does not plausibly extend to practical *foreknowledge*.³³

To see this, notice first that at least on the face of it, the idea of practical foreknowledge does seem even more puzzling than its contemporaneous parallel. Proponents of practical contemporaneous knowledge are challenged to explain why evidential support is not required for the belief that one is *V*-ing despite various possibilities that one's intention fails to exclude – e.g. that the window is not opening as one intends because there is a nail in the sash, or because the mechanism is faulty, or ... And there are so many *more* such possibilities when foreknowledge is at issue. These include scenarios of the sort used in the previous section to argue that one might not believe that one is going to *V* as one's intention may be abandoned prematurely. Even if, as Velleman suggests, one's belief is self-supporting; and even if, as Setiya suggests, one's belief is supported by one's knowledge-how – still, this does not guard against scenarios where the intention is extinguished before the time comes to act through forgetfulness, change of heart, dithering, and so on. Such scenarios seem (at least sometimes) to represent close possibilities that do not similarly threaten knowledge of current action. At least initially, then, practical foreknowledge

³² See, for example, Grice (1971), Langton (2004), and Paul (2009b).

³³ Moran (2004) also endorses foreknowledge of action.

seems considerably less plausible. Hence, any explanation of practical knowledge that ties its fortunes to practical foreknowledge should be considered *prima facie* less likely to succeed.³⁴

The intuitive expectation is confirmed when attention is turned to the resources that some non-cognitivist proponents of self-knowledge of action call on to explain the phenomenon. For these include important features of action ascriptions in the *progressive* (i.e. that one is *V*-ing) which do not similarly apply to the *prospective*. One such feature is the so-called ‘broadness’ of the progressive. Broadness consists in the independence of intentional action ascriptions from any particular bodily movements taking place in the narrow, more localized sense. It is illustrated when one picks an unusual route to get home: one is still correctly described as going home, even if at certain points along the way one is actually walking in the *opposite* direction from where one lives. The broadness of the progressive demonstrates that observation and inference are often not merely insufficient but strictly irrelevant to determining what one’s movements are intentionally bringing about: one’s intention to be going home overrides the observable bodily movements which suggest the contrary. Similarly, no amount of observation will reveal that one is making bread when one can be seen reading a newspaper on the couch. Yet one’s belief that one is making bread is not falsified by one’s reading; one may simply be waiting for the dough to rise.³⁵

Another feature peculiar to action ascriptions in the progressive is their so-called ‘openness’. Openness consists in the fact that events in general, and actions in particular, may be ongoing without ever having been completed. Thus one may be flying oneself to Nepal even though one will never actually make it to Nepal because as it happens, one is about to make a fatal error causing the plane to crash; similarly if one has a change of heart and switches destinations en route. As with broadness, openness arguably helps to explain why observation and inference are not required to support one’s belief that one is actually *V*-ing as one intends. For the original action-ascription is unconstrained by successful completion of one’s activity, and so neither is the corresponding knowledge-ascription. One need not wait to *see* that one has indeed *V*-ed in order to be warranted in believing that one is *V*-ing or was *V*-ing.

³⁴ The cognitivist may wish to exclude scenarios where the intention is extinguished prematurely by restricting practical foreknowledge to cases where such scenarios are unlikely to materialize. But absent a principled reason for doing so, the restriction would seem *ad hoc*.

³⁵ The examples are taken from Falvey (2000). See also Galton (1984: ch. 7, particularly pp. 122-8), and Thompson (2012).

In addition to interrupted actions and actions aborted by a change of mind, various possible errors in realizing one's intention likewise do not seem to falsify one's belief. Suppose one goes to the kitchen with the intention of making Breakfast tea, but absentmindedly reaches for a Chamomile teabag instead. Noticing the mistake does not convince one that it is actually Chamomile that one is making; one simply replaces the content of the cup with the right stuff. Overcoming a hitch is often enough itself a phase of the action, and does not force a retraction of one's original statement that one is *V*-ing. Unscrewing a bolt is not a sign that one is disassembling a machine contrary to one's intention, if one corrects the error and continues to assemble it.³⁶

The broadness and openness of progressive action-ascriptions arguably help to make plausible the idea of groundless self-knowledge of action by underscoring the relative *safety* of the belief that one is *V*-ing. Epistemic safety encodes an intuitive condition on knowledge whereby beliefs that amount to knowledge are not true thanks to accidentally favorable circumstances. An accidentally true belief would quickly degenerate into false belief if circumstances were to change slightly; knowing is modally more robust than that. And the openness and broadness of the progressive help to confirm the relative safety of one's belief that one is *V*-ing as one intends, in the face of the various possibilities that would seem to falsify it.

However, crucially these theoretical resources are unavailable when it comes to demonstrating that one's belief that one *will V* is likewise relatively safe. For they are resources generated by distinctive features of the progressive. This confirms the initial expectation above that a plausible account of practical knowledge, if there is one, would not extend to cover foreknowledge – unlike the account offered by the cognitivist. Positing a belief in prospective intention as WIC does turns out to be an unwarranted overgeneralization in the present context, since the phenomenon of practical knowledge is at best restricted to the progressive.

Notice that the problem cannot be avoided by appealing to how detailed the content of the belief is (as an anonymous referee suggested it might). Plausibly, other things equal, the safety of one's belief

³⁶ There are plausibly limits to how gross the errors could be for it still to be the case that one is *V*-ing. Furthermore, Falvey (2000) concedes that observation here does play the auxiliary role of revealing to one that adjustments are required to bring the action back on course. But it is not the method by which one knows what one is intentionally doing. As Anscombe notes (1957: 53), with eyes closed one may perhaps fail to write *legibly*; but one is still writing.

decreases the further in the future one's V -ing is supposed to take place, as there are more close possibilities that one will not V . But equally plausibly, safety increases the sparser the content of the belief, i.e. the less specific V -ing is. For example, other things equal, the belief that one will have tea is safer than the belief that one will have chamomile tea. And one might suggest that the latter increase can compensate for the former decrease and restore sufficient safety: intentions that are further in the future, and hence their corresponding beliefs, tend to be sparser in content since they do not yet require detailed planning. However, for one thing, we do seem to have fairly detailed intentions for the remote future – e.g. the intention to meet up with our high school friends at the annual reunion in the school auditorium. Furthermore, however sparse the intention is, it will remain vulnerable to scenarios mentioned above, in which it is extinguished before the time comes to execute it, due to change of heart, forgetfulness, and so on. (Indeed, other things equal, the likelihood of extinction surely rises the further in the future the intended action takes place).

To summarize the discussion in this section so far, the argument from practical knowledge overgeneralizes, implausibly implying the existence of practical foreknowledge. WIC can therefore hardly be considered the best explanation of practical knowledge. And consequently, WIC cannot claim support from how it best explains this phenomenon.

A possible response circumscribes WIC accordingly. If the idea of practical foreknowledge is an undesirable implication, perhaps WIC should be more narrowly construed as applying to intention-in-action alone. Indeed, assuming that knowledge entails belief, the considerations above from the broadness and openness of the progressive may be taken to support the view that in acting with the intention to V , one knows – and hence *believes* – that one is V -ing. However, cognitivists would be ill-advised to adopt such a restriction. For one thing, it would undercut an important advertised feature of the view. The feature in question is the identification of the *unifying* mark of intention across its different manifestations, including prospective intention and intention-in-action.³⁸ Furthermore, the restriction to intention-in-action would similarly undermine some of the salient theoretical payoffs of WIC described in section I,

³⁸ Cf. Setiya (2015): '[T]he basic thought is that intention in action involves the belief that one is doing A . Doing something for a reason involves a belief about one's reason for doing it that constitutes intention in action. And prospective intention, or intention for the future, involves a belief about what one is *going* to do and why. *Intention as involving belief is the thread that binds these phenomena together.*' (Emphasis added). And see also Setiya (2007b: 48-9). On the need to explain how the different manifestations of intention are related, see Anscombe (1957: 1).

such as the cognitivist explanations of instrumental rationality and intention consistency, and the possibility of a unified explanation of practical and theoretical reasoning, seeing as the former type of reasoning typically concludes in a *prospective* intention.³⁹ And lastly, restricting the scope of the thesis in this way seems objectionably *ad hoc*. After all, the worries motivating the restriction do not stem from general grounds for doubting that prospective intention in particular entails belief, which may have licensed a restriction to intention-in-action. Rather, the worries stem from flaws in one of the arguments offered as support for the thesis.

If the extension to prospective intention is non-negotiable, the cognitivist may instead attempt to defend its plausibility in the face of the above objections. One way of doing so would be to try to identify other beliefs that *do* plausibly amount to knowledge and display a similar epistemic profile to that of beliefs that one will \checkmark as one intends. For example, we ordinarily take ourselves to have *non*-practical foreknowledge of some propositions. Thus I seem to know that my spouse will be home around 7pm as she said she would. Nevertheless it is possible that her car will be stolen, forcing her to go to the police station instead. The example arguably shows that to have knowledge of future states and events, not all possible scenarios that would falsify what is known need be ruled out. Hence it might be claimed that the above objection to practical foreknowledge proves too much, insofar as it sets a constraint that is not met by perfectly credible instances of non-practical foreknowledge.

Putting external world scepticism to one side, it would no doubt be highly implausible to suggest that knowledge in general, and knowledge of action in particular, requires that *no* contrary possibilities exist. However, assuming knowledge of action must be safe, it does require that no *close* contrary possibilities exist. And the sort of contrary possibilities highlighted above – those caused by one changing one’s mind, forgetting, dithering, etc. – might well be, at least on *some* occasions, quite close. If you are heading to the cafeteria, intending to have soup for lunch, then plausibly you need not rule out forgetting or abandoning this intention to know that you will indeed have soup. But if you intend to visit that quaint little café when you are next in Paris, then plausibly you do. And as noted above (n. 34), circumscribing practical knowledge to cases where such possibilities are remote seems *ad hoc*.

³⁹ I’m indebted to an anonymous referee for pointing out the upshot of a disunified account of reasoning.

A different attempt to render plausible the application of WIC to prospective intention reformulates the view in terms of partial belief or degrees of confidence or credence, rather than outright belief. For example, perhaps intending to V (occasionally?) entails merely being more confident that one will V than one would otherwise be.⁴⁰ Modified in this way, the view will avoid the problematic implication of practical foreknowledge that one will V . For, at least in the problematic cases where it seems one cannot have such knowledge, one may not in fact (outright) believe that one will V . The resulting view is a watered-down, far less striking or controversial version of WIC: in intending to V , one may have a rather slight degree of confidence that one will V , so long as it is higher than one would have had absent the intention. More importantly, however, and much like the possible restriction to intention-in-action considered above, this move is a double-edged sword. For it strips WIC of important advertised virtues of the view, including the claim that intention-as-belief is the unifying mark of intention; that belief is the conclusion of practical (and not just theoretical) reasoning; and that practical rationality is a species of theoretical rationality. All these require that intention entail outright belief.

IV

This concludes the paper's attack on Intention Cognitivism. Several doubts concerning the motivations for the view have been raised. Section II argued that, contrary to what proponents of IC suggest, the fact that intentions and beliefs are standardly expressed in similar ways does not support the view. For, first, as argued in II.1, it is unclear that intentions are actually expressed by asserting propositions rather than by directly giving voice to one's intention, as avowal-expressivism maintains. And if expressivism delivers the correct way to understand expressions of intention, the latter need have nothing distinctive in common with expressions of belief. Second, as argued in II.2, even if intentions are expressed by asserting, so that there *is* a uniquely salient parity of expression with beliefs, this fact may be explained by how expressing an intention conversationally implicates the presence of belief, rather than entailing it. The remainder of section II was devoted to questioning, on similar grounds to those invoked in II.2, a further possible

⁴⁰ Thanks to an anonymous referee for urging me to consider this response. Setiya (2009: 130) suggests such an adjustment of WIC to partial belief in response to a different worry, viz. that an agent who is V -ing intentionally may in some circumstances fail to believe that she is V -ing.

motivation for IC: the cognitivist's claim that the possibility of planning around one's intention – either inter- or intra-personally – presupposes the truth of WIC. And finally, section III argued that WIC does not offer a plausible vindicating explanation of 'practical knowledge'. For that explanation overgeneralizes and distorts the phenomenon.

The overarching point of raising these objections to IC has been to attempt a systematic assessment of the grounds taken to motivate the view in the first place. And the outcome of this assessment is unfavourable to IC; all the motivations surveyed were found wanting. As such, the cumulative upshot of the foregoing discussion falls short of conclusively *refuting* IC. But the discussion has exposed IC – and its potentially significant payoffs spelt out in section I – as unfounded. Hence it suggests we should withhold judgment on WIC, pending further, more compelling grounds to endorse the view⁴¹.

References

- Anscombe, Elizabeth 1957. *Intention*, Basil Blackwell.
- Bar-On, Dorit 2004. *Speaking My Mind: Expression and Self-Knowledge*, Oxford University Press.
- Bar-On, Dorit, and Long, Douglas 2001. Avowals and First-Person Privilege, *Philosophy and Phenomenological Research* 62: 311-335.
- Bar-On, Dorit and Sias, James 2013. Varieties of Expressivism, *Philosophy Compass* 8: 699-713.
- Bratman, Michael 2009. Intention, Belief, and Practical Rationality, in *Reasons for Action*, eds. D. Sobel and S. Wall, Cambridge University Press: 13–36.
- Bratman, Michael 1987. *Intention, Plans, and Practical Reason*, Harvard University Press.
- Broome, John 2013. *Rationality Through Reasoning*, Wiley-Blackwell.
- Broome, John 2005. Does Rationality Give Us Reasons? *Philosophical Issues*, 15, 321–337.
- Brunero, John 2009. Against Cognitivism about Practical Rationality, *Philosophical Studies* 146: 311-325.
- Davidson, Donald 1980. *Essays on Actions and Events*, Oxford University Press.
- Donnellan, Keith 1963. Knowing What I am Doing, *Journal of Philosophy* 60: 401-409.
- Falvey, Kevin 2000. Knowledge in Intention, *Philosophical Studies* 99: 21-44.
- Galton, Antony 1984. *The Logic of Aspect: an Axiomatic Approach*, Oxford: Clarendon Press.

⁴¹ For their invaluable written comments on earlier drafts of this paper and discussion of material appearing therein, I'm extremely grateful to John Broome, Alex Gregory, John Hyman, Conor McHugh, Noa Leibowitz, Kurt Sylvan, Lee Walters, Jonathan Way, Daniel Whiting, audiences at Oxford University's Moral Philosophy Seminar and the Southern Normativity Group's 2014 conference at Kent, and two anonymous referees. Thanks also to the UK's Arts and Humanities Research Council for funding in support of this research (grant no. AH/K008188/1).

- Green, Mitchell, and Williams, John (eds.) 2007. *Moore's Paradox: New Essays on Belief, Rationality, and the First Person*, Oxford University Press.
- Grice, H. P. 1989. *Studies in the Ways of Words*, Harvard University Press.
- Grice, H. P. 1971. Intention and Uncertainty, *Proceedings of the British Academy* 5: 263–279.
- Harman, Gilbert 1976. Practical Reasoning, *Review of Metaphysics* 29: 431-463.
- Heal, Jane 1994. Moore's Paradox: A Wittgensteinian Approach, *Mind* 103: 5-24.
- Holton, Richard (2008). Partial Intention, Partial Belief, *Mind* 117: 27-57.
- Kolodny, Nico 2005. Why be rational? *Mind*, 114: 509-563.
- Langton, Rey 2004. Intention as Faith, in *Agency and Action*, ed. J. Hyman and H. Steward, Cambridge University Press: 243–258.
- Moran, Richard 2004. Anscombe on 'Practical Knowledge', in *Agency and Action*, eds. J. Hyman and H. Steward, Cambridge University Press: 43-68.
- Paul, Sarah 2009a. Intention, Belief, and Wishful Thinking: Setiya on 'Practical Knowledge,' *Ethics* 119: 546–57.
- Paul, Sarah 2009b. How We Know What We're Doing, *Philosophers' Imprint* 9: 1-24.
- Raz, Joseph 2005. The Myth of Instrumental Reason, *Journal of Ethics and Social Philosophy* 1: 1-28.
- Ridge, Michael 1998. Humean Intentions, *American Philosophical Quarterly* 35: 157-178.
- Ross, Jacob 2009. How to be a Cognitivist about Practical Reason, *Oxford Studies in Metaethics* 4: 243–281.
- Setiya, Kieran 2015. Intention, in *The Stanford Encyclopedia of Philosophy*, ed. E. N. Zalta, URL = <http://plato.stanford.edu/archives/sum2015/entries/intention/>.
- Setiya, Kieran 2011. Knowledge of Intention, in *Essays on Anscombe's 'Intention'*, ed. A. Ford, J. Hornsby and F. Stoutland, Harvard University Press: 170-197.
- Setiya, Kieran 2009. Practical Knowledge Revisited, *Ethics* 120: 128–137.
- Setiya, Kieran 2008. Practical Knowledge, *Ethics* 118: 388–409.
- Setiya, Kieran 2007a. Cognitivism about Instrumental Reason, *Ethics* 117: 647–673.
- Setiya, Kieran 2007b. *Reasons without Rationalism*, Princeton University Press.
- Sinhababu, Neil 2013. The Desire-Belief Account of Intention Explains Everything, *Noûs* 47: 680-696.
- Thompson, Michael 2012. Anscombe's 'Intention' and Practical Knowledge, in *Essays on Anscombe's 'Intention'*, eds. A. Ford, J. Hornsby, and F. Stoutland, Harvard University Press: 198-210.
- Velleman, David 2007. What Good is a Will?, in *Action in Context*, eds. A. Leist and H. Baumann, de Gruyter/Mouton: 193-215.
- Velleman, David 1989. *Practical Reflection*, Princeton University Press.
- Wright, Crispin 1998. Self-Knowledge: The Wittgensteinian Legacy, in *Knowing Our Own Minds*, ed. C. Wright, B. Smith, and C. Macdonald, Oxford: Clarendon Press: 13-45.