

Intending is Believing: A Defense of Strong Cognitivism

Berislav Marušić, Brandeis University
John Schwenkler, Florida State University*

We argue that intentions are beliefs—beliefs that are held in light of, and made rational by, practical reasoning. To intend to do something is neither more nor less than to believe, on the basis of one’s practical reasoning, that one will do it. The identification of the mental state of intention with the mental state of belief is what we call *strong cognitivism* about intentions. It is a strong form of cognitivism because we identify intentions with beliefs, rather than maintaining that beliefs are entailed by intentions or are components of them.

To illustrate the basic idea, consider the attitude that you might express with the words, “I am going to fail this exam” (Anscombe 1957/2000, 2). We propose that this attitude is a practical conclusion if it is grounded in practical reasoning—if you hold it because you see failing the exam as somehow worthwhile. (Perhaps it will impress your love interest.) It is a theoretical conclusion if it is held in light of evidence. (Perhaps you think that you have not studied enough to do well.) We hold that both conclusions constitute beliefs and exhibit the functional roles of belief. The former is an intention or, as we shall sometimes say, a “practical” belief, the latter an evidence-based, “theoretical” belief.

The aim of strong cognitivism is to explain the relation between practical and theoretical reasoning—that is, to explain how reasoning about *what to do* relates to reasoning about *what will happen* in the future. This relation is difficult to understand if practical and theoretical reasoning are seen as answering different sorts of questions—if, say, the conclusion of practical reasoning is supposed to be a judgment about what it would be good to do, or what one ought to do, or what one has most reason to do. By contrast, we hold that these two forms of reasoning are related because practical reasoning is concerned with a factual question—the question of *what one is going to do*. This is why we say that practical reasoning and theoretical reasoning both issue in beliefs.¹ What distinguishes the conclusions of practical and theoretical reasoning is not what they are about but how they are grounded—that is, the form of reasoning in light of which the respective conclusions are held.

In what follows, we first describe what we take to be the essential features of practical reasoning and explain how practical reasoning can ground belief (§1). We then offer a series of arguments for strong cognitivism (§2) and address a number of objections to it (§3). We conclude with some general remarks about why strong cognitivism of the kind we defend is philosophically important.

* Both authors contributed equally to this work.

¹ In saying that these forms of reasoning issue in belief we don’t mean to deny that often (or in the good case) they issue in something more: knowledge of what will happen, or intentional action of a given kind. But we take it for granted that knowledge and action both require belief.

1. Practical Reasoning

It is beyond the scope of this paper to offer a detailed theory of the distinction between practical and theoretical reasoning. However, since this distinction is so important to our position, we outline below how we propose to understand it (§1.1). We then explain how it is possible for practical reasoning to issue in belief (§1.2).

1.1 Distinguishing Features of Practical Reasoning

We propose that practical reasoning is the reasoning by which an agent determines what she is going to do, insofar as what she will do is up to her. That is, insofar as matters are up to the agent, she is in a position to determine how they will unfold, and she does this through her practical reasoning. Since an agent's actions are grounded in the reasoning by which she determines what she will do, the truth she creates by acting is not independent of her reasoning about it. This means that an agent's view of what will happen does not simply *reflect* the future as it will be anyway, independent of her reasoning about it, but also helps to *create* the future it represents.² In deciding what to do through her practical reasoning, an agent determines what will be true in light of what she sees as good.

There is nothing mysterious in the idea that agents have a special way of forming beliefs about the future. Rather, it simply reflects the fact that insofar as a person's actions are under her control, her own beliefs are part of what determines what she will do. This means that from an agent's perspective, the question "Will I ϕ ?", where whether or not she ϕ s is up to her, *has* the answer she *gives* it. There is no independent truth for her answer to reflect, and no way for the evidence to settle what this answer should be. That is why the agent has to address this question through practical reasoning, rather than through reasoning whose conclusion is meant to reflect what is independently likely to happen.³ For an agent's beliefs about what she will do *cannot* be grounded just in evidence about what will anyway happen—if matters are up to the agent, then this evidence will necessarily be inconclusive.⁴

To illustrate, return to the earlier example of believing that you are going to fail an exam, and contrast two ways that you might hold such a conclusion. First, the conclusion might be held in light of considerations that make it appear *likely* that you will fail—that is, considerations that make this seem *independently* likely, such that you are not in a position to decide whether or not you will pass. Second, the conclusion might be held in light of considerations that make failing the exam appear *worthwhile* or *good*. In this second case, the truth about how you will do on the exam is *not independent* of the conclusion you reach—for this matter is up to you to decide. In the first case your reasoning is theoretical and your conclusion is a theoretical belief or, as G.E.M.

² The terminology of "reflecting" and "creating" truth is due to Velleman (1996, 195, n.55). More on this shortly (§1.2).

³ Of course it is possible to *try* to form such beliefs about one's future in the latter way. We hold that doing so is a form of what Sartre called bad faith. For discussion, see Marušić (2013; 2015, ch. 5.3).

⁴ This is a central theme in Bok (1998). See also Ismael (2012).

Anscombe calls it, an “estimate of the future” (1957/2000, 3). In the second case, your reasoning is practical and your conclusion is a practical belief, which we identify with an intention.

In characterizing the nature of practical reasoning in this way we do not deny that considerations of what is independently likely can figure in practical thought, nor that considerations of goodness or worth can figure in reasoning about what is independently true. The difference we aim to characterize is in the *form* of reasoning rather than the *content* of the considerations related in it (Anscombe 1957/2000, 60); it concerns the *manner* in which a conclusion relates to its grounds rather than the *matters* that the conclusion or its grounds concern. We propose that the form of practical reasoning is essentially instrumental or *calculative*: Through practical reasoning, an agent concludes that she will do something because she regards this action as a suitable means to an end, or sometimes as a suitable end in itself. That is, in practical reasoning the conclusion that one will do something is held in light of the appearance that such an action would be a suitable way to do, acquire, or become something that one desires. In offering this characterization we do not deny that it is possible also for reasoning to determine the desires in light of which we decide to act.⁵ Our claim is that there is *practical* reasoning only where there is reasoning with this distinctively instrumental or calculative form—reasoning through which a person determines *how* she is going act, or *what* she will do in order to achieve her ends.⁶ This is what makes practical reasoning different from reasoning that proceeds in light of evidence for the truth of its conclusion.

Following Anscombe (1957/2000, 5ff), we can also mark the difference between practical and theoretical reasoning by distinguishing two senses of the question “Why?” that can find application in reference to an action. In the case that interests us, where the action in question is in the future, the question “Why are you going to do that?” asks for an agent’s practical reasoning when it asks her to explain the *point* of what she proposes to do, or the considerations in light of which her action appears as worthwhile or good. By contrast, the question asks for her theoretical reasoning when it asks her to explain the factors that make it seem independently likely that she will do this, and thus to identify the probable *causes* of the behavior that she estimates will happen. As we will emphasize below, the very same considerations could be cited in an explanation of either sort: One might, for example, either *choose* to fail a test because of its apparent difficulty (to prove a point, perhaps), or *estimate* that she will fail it for this very same reason. What makes the first case an instance of practical reasoning, and the second theoretical, is the *way* that the conclusion—that one will fail the exam—is drawn from its ground in each case: In reasoning practically one supposes that the difficulty of the test shows failing it to be good or worthwhile, whereas in theoretical reasoning one supposes this to show that failure is likely. We hold that a belief that one will do something is an intention when and only when it grounded in reasoning of the first sort—that is, a belief is an intention when the question “Why?”, where this asks for the point or purpose of a proposed action rather than its causes, can be given application in regard to it.

⁵ For discussion, see Korsgaard (1986).

⁶ The calculative or “how-to” structure of practical reasoning is explored, in different but complementary ways, by Vogler (2002) and Ford (2017).

This way of distinguishing practical and theoretical reasoning contrasts with some other influential accounts of this distinction. We reject the suggestion that practical reasoning is distinctive because it issues in intention rather than belief (Broome 2002), since on our view an intention is just a belief of a special sort—namely, one grounded in practical reasoning. We also reject the view that practical reasoning is ordinary reasoning “when applied to practical matters” (Raz 2011, 129), such as the question of what it would be good to do, or what one ought to do, or what one has most reason to do. We hold, with Anscombe (1957/2000, 58), that reasoning that addresses such questions is no more a distinctive *kind* of reasoning than is reasoning about mince pies: What makes practical reasoning practical is something other than what it is about, as it is possible to determine what is good, obligatory, favored by reasons, and so on in a wholly theoretical way, drawing a conclusion that is meant to reflect an independent normative landscape. Indeed, on our view there *is* nothing distinctive about the conclusion of practical reasoning, taken in itself: The very same conclusion, namely that one will do such-and-such, that is reached through practical reasoning can also be reached theoretically, as something one estimates will happen in light of evidence about what one will do. On our view, what makes practical reasoning practical is the way that its conclusion is supposed to be supported by the considerations in light of which one holds it: The conclusion of practical reasoning is not held because it is “shewn to be true by the premises” (Anscombe 1957/2000, 58); rather, it is held because its truth is shown by the premises to be good.

The difference between practical and theoretical reasoning is reflected in the distinctive way that practical beliefs are responsive to evidence of an action’s difficulty. For an agent, considerations showing that a certain action will be difficult need not lead her to conclude that she will not act in this way, since if the action is up to her and important enough, the rational response to evidence of difficulty may be to find ways to overcome it, say by avoiding temptation, strengthening resolve, making strategic choices, and so on. At other times this may not be rational: Evidence of difficulty, and of what it would take for the difficulty to be overcome, might show the action not worthwhile after all, because it is not worth the effort. In this case, too, one’s reasoning would be distinctly practical.⁷ But things are different for beliefs that are supposed to reflect what is independently the case: A person who believes in light of evidence that she will do something, and then accepts evidence that she won’t do this after all, is rationally obliged to revise her belief. We suggest that the explanation of this difference is that predictive beliefs are grounded in theoretical reasoning, whereas practical beliefs, or intentions, are grounded in practical reasoning, and so do not aim to reflect what is anyway going to happen.⁸

Let us stress again that in claiming that intentions are beliefs that are grounded in practical reasoning, we do not mean to preclude that they are sensitive to evidential

⁷ See Marušić (2012; 2015, ch. 6.1) and also Holton (2009, 75).

⁸ It is only when the evidence suggests that a given action is *impossible* for a person that she can’t respond to it with an effort to surmount the difficulty, but is rationally required to weaken or revise her belief that she will act as she intends—a form of weakening or revision we discuss in detail in §3.1 below.

considerations. For example, when Jack is considering whether to watch a movie in the theater, he will be concerned not only with whether it is worthwhile to see the movie, but also with whether the movie is, in fact, playing at the theater and whether he can get to the theater on time. What makes Jack's conclusion to see the movie an intention is that he relies on this evidence, not to *predict* what he is going to do, but to *decide* this, in light of his view of whether or not seeing the movie is worthwhile. Evidence of when the movie will be playing is relevant to this decision because this evidence is relevant to determining whether, and when, it is worthwhile to go and see it. That he consider evidence in this way does not render his reasoning theoretical.

Finally, we hasten to add that something can be the conclusion of practical reasoning, and hence an intention, even if there was no *conscious* or *explicit* deliberation leading up to it. What is criterial of a conclusion of reasoning is the sense of the question "Why?", where this seeks the reasons *in light of which one holds it*, that can be given application to it. For a predictive belief or "estimate of the future" the applicable sense of this question is one in which it asks for considerations that make it appear likely that one will do something, whereas for an intention or practical belief the applicable "Why?"-question asks for the considerations in light of which the action appears as good or worthwhile—as either an end in itself or a suitable means to one or more of the agent's ends. There is room to make these distinctions even if the believer didn't reason explicitly, even if she doesn't have a particularly revealing explanation to offer (for instance, if she answers, "For no reason"), and even if she offers an outright bad explanation—if she has reached her conclusion in light of what are in fact, and possibly even by her own lights, bad reasons. None of this precludes the possibility of distinguishing these two ways in which the question "Why?" can be posed and given application—one way in which it asks for an evidential basis in light of which the proposed action can be predicted, the other in which it asks for a practical basis in light of which the action appears as good.

This concludes our summary account of the nature of practical reasoning. Certainly there is much more to be said on this subject. But we hope that this is enough to clarify our position, which will be developed in more detail below.

1.2 Belief as the Conclusion of Practical Reasoning

The thesis that intentions are beliefs grounded in practical reasoning may strike many readers as a non-starter since, as it is widely held, belief is—in an important but hard-to-articulate way—essentially related to truth. In light of this worry, we set out in this section to address a conceptual question:⁹ How *could* intentions be beliefs if, to borrow Williams's famous slogan, "belief aims at truth" (1973)? In particular, how could intentions be beliefs, since belief aims at truth—and, some say, this just means that belief must aim to reflect evidence about what is or will be the case (Adler 2002; Shah 2006)? We claim that intentions are beliefs formed in light of practical reasoning—but it is an important part of our view that practical reasoning does not simply reflect evidence about what is independently the case. Thus it may seem that such reasoning can only support

⁹ See also Marušić (2015, ch. 6.2).

the conclusion, not that something is or will be *true*, but rather that something is *worthwhile*.

To answer this objection, we first have to clarify how practical reasoning can be the ground of a belief that one will φ . On our view, practical reasoning does this by showing φ -ing to be worthwhile, and not by showing it worthwhile to *intend* to φ or *believe* that one will do so. That is, we hold that it is only in light of finding it worthwhile *to* φ that a person can rationally come to intend to φ and hence to believe that she will φ . It is only in this way that practical reasoning can make intending—and hence believing—rational. Thus we do not countenance—indeed we reject—pragmatism in the ethics of belief, according to which the fact that it would be worthwhile *to believe* something can be a good reason to believe it.

With this preliminary remark in hand, we can turn to the conceptual question: How could intentions be beliefs that aim at truth? Our response is to distinguish *truth-commitment*, which is a characteristic of belief in general, from *being held in light of evidence*, which is a characteristic only of some beliefs, namely theoretical beliefs. To see the difference between these notions, suppose for the sake of argument that, as David Velleman puts it, “to believe something is to accept it with the aim of doing so only if it is really true” (1996, 185)—that is, to accept it in a way such that one is *committed* to its truth. (We could add, following Shah (2003), that to believe is furthermore to take truth to be the standard of correctness for one’s acceptance. We will omit this for ease of exposition, but it can be read as implicit throughout our argument.)¹⁰ As Velleman points out, there are then two forms of such acceptance: “accepting [it] so as to reflect the truth” and “accepting [it] so as to create the truth” (1996, 195, n.55). But only the first form of acceptance requires that one’s attitude be held in light of evidence about what is going to happen. One may also hold a *practical* belief, therein accepting something so as to *create* truth, not in light of evidence that this is going to happen, but on the basis of considerations that show this truth to be worth creating. Our position is that intentions are beliefs of this second sort: An intention is a commitment to *make* something true, in light of the fact that this truth appears worth creating. Thus we hold that practical reasoning *does* support the conclusion that something is (or will be) true. What distinguishes practical from theoretical reasoning is that practical reasoning supports such a conclusion *in light of the fact that this truth appears as good*. Intentions are truth-commitments. But they are the truth-commitments of an agent—they are truth-commitments grounded in practical reasoning.¹¹

This completes our reply to the conceptual question. The truth-commitment of intentions is secured by the fact that intentions are the commitments of an agent. And it is the prerogative of an agent to settle the question of what will be true, insofar as this is up

¹⁰ We hope this makes clear that our argument is compatible with other views of what it means for belief to aim at truth, such as Wedgwood’s (2002) and also Velleman’s later view defended together with Shah (Velleman and Shah 2005).

¹¹ Doyle (MS) makes a related point, writing that “As long as the point of my description of the future is *to be true*, whether in the one way, as expressing my best estimate of how future events will unfold, while I (so to say) *look on*, or in the other, as something to be *made* true by my own action, there is warrant for calling that description the content of my *belief* ...”

to her, in light of what appears to her as worthwhile or good.¹²

The preceding discussion also suggests a simple argument in favor of the strong cognitivism that we defend. If beliefs are essentially truth-commitments, then there must be a way of believing that arises from an agent's capacity to create the truth: There must be characteristically agential truth-commitments. If there are beliefs of this sort—practical beliefs—then we must be able to find them in our mental economy. Strong cognitivism provides a straightforward way of doing this, as it holds that our practical beliefs just are our intentions.

This argument explains why we defend a form of strong cognitivism that identifies beliefs and intentions outright, rather than the weaker view that intentions merely have beliefs as constituents, or entail belief in success (Grice 1971; Harman 1976; 1986). Views that don't identify intentions with practical beliefs must distinguish theoretical beliefs and practical beliefs, and both in turn from intentions—a series of distinctions that introduces complications where none are needed.¹³ Once we are clear about the difference between practical beliefs and theoretical beliefs, we would need to find a principled reason to distinguish practical beliefs from intentions. We consider some candidate reasons below in Section 3, arguing that none are sufficient to warrant distinguishing intention from belief. Before this, in Section 2 we offer some further arguments for our position.

2. Arguments for Strong Cognitivism

Our three arguments in favor of strong cognitivism are the following: Strong cognitivism explains why beliefs about what we are going to do are usually transparent to practical reasoning (§2.1); it is supported by considerations about the expression of intentions (§2.2); and it explains why intentions serve a belief-like role in planning (§2.3). As we emphasize below, these arguments are not intended to provide definitive support for our view over the alternatives to it. Rather, our overall argument has the form of an inference to the best explanation: We hold that strong cognitivism provides the best—because most elegant, intuitive, and unforced—explanation of the phenomena we discuss, and should be favored for that reason.

2.1 Transparency

Our first argument in favor of strong cognitivism is that it offers the best explanation of why our beliefs about our future actions are usually settled in light of practical reasoning. That is, it explains why, in ordinary circumstances, a person will address the

¹² Our view thus does not extend to self-fulfilling beliefs—beliefs that come true because one holds them. Intentions are, on our view, not leaps of faith (Langton 2004). Our point is rather that intentions are beliefs because they are the beliefs of someone who is committed to making these beliefs true.

¹³ Hieronymi (2009, n.3) explicitly distinguishes intentions from predictions on the basis of practical reasons, which, in turn, are distinct from predictions on the basis of theoretical reasons.

question of whether she will φ through practical deliberation about *whether to φ* , and thus by considering whether or not φ -ing is worthwhile, rather than weighing evidence for and against the judgment that she will φ .

To see that this is the case, imagine that Daniel is considering whether or not he will go to a wedding that he's been invited to, and which he clearly has the ability to attend. Answering the question of whether he will attend the wedding seems to require having a belief about the matter, since the question is about whether a certain thing will happen in the future. (Perhaps he is considering the question because a friend has asked whether he will be there, as she wants to see about sharing a hotel room.) Yet in the ordinary case at least, the way Daniel will form such a belief is not by considering evidence of the likelihood that he will or won't go, but rather by *practical* reflection of the sort that issues in an intention or decision: In arriving at an answer to this question Daniel will consider e.g. whether the wedding will be enjoyable, whether he is under any obligation to attend, how expensive it will be to get there, whether the date of the wedding conflicts with other things he intends to do, and so on. Following Richard Moran (2001), we will put this point by saying that Daniel treats the question of whether he will go the wedding as *transparent* to the question of whether *to go*—a question that he settles in light of considerations that reveal the goodness or worth of doing this.¹⁴

Though the transparency we are describing here seems to be an ordinary phenomenon, it may be puzzling how that phenomenon could be anything more than an illusion. For isn't the question of whether one *will do* something related only indirectly to considerations bearing on whether doing that thing is good or worthwhile? At the very least, there seems to be an inferential step between deciding to do something and concluding that I will do it (Paul 2009): It seems that a person would have to infer, on the basis of having decided to do something, that she will actually do it. And the goodness of this inference would depend on her evidence about how effective her decision is likely to be.¹⁵ Given this, why is it not simply irrational to treat the question of what one will do as one that can be resolved through practical deliberation? Why is this not a flagrant disregard of evidence concerning the effectiveness of our decisions?

Our version of strong cognitivism provides an explanation of this phenomenon that avoids this appearance of irrationality. We propose that it is because the intention to φ that is grounded in practical reasoning simply *is* a belief that you will φ , that the question of whether you will φ can be answered through practical deliberation. That is, we hold that this is possible because your answer to the question of whether to φ , which you reach through practical deliberation in light of considerations pertaining to the goodness or worth of φ -ing, simply *is* your belief about what you are going to do. It is because of the identity of belief and intention that it can be rational to answer the question of what you *will* do in the very same way as you form the intention *to* do this.

¹⁴ Moran holds that this transparency is a "kind of normative ideal" (2001, 62), and similarly we say that it holds usually, or in the ordinary case. Thus our position, like Moran's, is explicitly meant to accommodate cases where the deliberator might be irresolute or open to changing her mind, or where there is a concern that circumstances might make it impossible for her to do what she intends. We will consider in §3.1 whether these cases are counterexamples to our strong cognitivism.

¹⁵ See Wilson (2000) for trenchant criticism of this sort of view.

In offering this argument we do not deny that this phenomenon of transparency could be explained in other ways than by appeal to our strong cognitivism. For example, it might be that the question of what one will do can be answered by practical reasoning because intention and belief are linked in a relation other than that of identity, such as through inference or a rational requirement to believe that one will act as one intends, or because intention has belief as a constituent part. Our argument is that strong cognitivism offers the *best* explanation of this phenomenon, insofar as the explanation does not rely on any further assumptions: The most straightforward explanation of why the question of what one will do is answered through practical reasoning is that the conclusion of practical reasoning *simply is* one's answer to this question.

2.2 Expression of Intention

Our second argument for strong cognitivism is that it is supported by consideration of the ways in which intentions are expressed.¹⁶ Here we can begin with the familiar idea that assertion expresses belief.¹⁷ Assertion is the outward, communicable side of belief. It is, as it were, the face of belief. As Williamson puts it, "It is plausible ... that occurrently believing *p* stands to asserting *p* as the inner stands to the outer" (2000, p. 255). That is why conclusions about the speech act of assertion will reveal something about the nature of the mental state of belief.

We hold that *by the same token*, conclusions about the speech act of assertion reveal something about the nature of the mental state of intention. In particular, we hold that the fact that the usual expression of an intention is an assertion about what one will do, which is also the usual expression of a corresponding belief, is evidence that intention and belief are one and the same kind of state. That intention is expressed through assertion is a well-known point; it was noted by Anscombe, who writes at the opening of *Intention* that "Very often, when a man says 'I am going to do such-and-such', we should say that this was an expression of intention" (1957/2000, 1).¹⁸ Strong cognitivism offers a straightforward explanation of why intention would be expressed in this way: Since outright assertion is an expression of belief, and since an intention is a belief about what one will do, the way to express an intention is to say what one will do. Since the same speech act can be used to express belief and intention, belief and intention share a face—and this is evidence that they are one and the same mental state.¹⁹

¹⁶ This section builds on arguments presented in Marušić (2015, ch. 3.1). For similar analysis, appealing to Anscombe's (1957/2000) discussion of the expression of intention, see Doyle (MS).

¹⁷ See Williams (1973, 137–8; 2002, 74), Bach and Harnish (1979, 47), Williamson (2000, 255), and Adler (2002, 13–14 and *passim*).

¹⁸ The point is also made, among others, by Williams (1973, 137) and, more recently, Setiya (2008, 396).

¹⁹ Davidson has famously argued that "a man who sincerely says, ... 'I will do it' under certain conditions must believe that he will do it. But it may be the saying, not the intention, that implies belief" (1978, 91). In arguing this, Davidson effectively denies that assertion expresses an intention: His point is that it merely expresses belief, since belief is required to properly assert, "I will do it." Davidson's view therefore leaves it unclear how one would express one's intention. What does one say, to oneself or others, to express one's *intention*, rather than one's belief? On a

Yet this argument leaves us with an open question: If outright assertion expresses an intention, one might wonder what we express with the phrase, “I intend.” We will return to this issue in detail when we consider the objection that we can intend to do something while doubting that we will succeed. For now, we just want to note a parallel between uses of the phrase “I intend” and uses of the phrase “I believe.” An assertion of the form “I believe that p ” typically expresses what we will call a *partial* belief, i.e. a conditional belief or a high degree of belief that falls short of outright conviction.²⁰ For example, when Marko says, “I believe that it is raining,” he conveys that he has a high degree of belief that it is raining but is not entirely sure. By contrast, when he says, “It is raining,” he conveys outright belief. This could be explained in Gricean terms (Grice 1989): Since Marko uses the phrase “I believe,” it may be presumed that he is not prepared to assert p outright. Since outright assertion expresses outright belief, choosing this form of less-than-outright assertion implies that he does not outright believe that p .

We hold that the situation is analogous with intentions.²¹ When Marko says, “I intend to go to Peter’s party,” he typically expresses a *partial intention*. (This concept is developed in more detail in §3.1 below.) In saying this, he conveys that he is to some extent committed to going but is not entirely sure that he will go—be it because his intention is conditional on certain facts obtaining, because he is unsure about whether he will succeed in going, or because he simply has not entirely made up his mind. He thus leave open the possibility that he will decide not to go, or that he will be somehow unable to get there. By contrast, when Marko says, “I will go to Peter’s party,” he conveys that the question of whether he will go is settled for him. In light of this, we hold that whatever is to be said about our use of the phrase “I believe” can also be said about our use of the phrase “I intend.”²² However, just as not all theoretical beliefs are expressed with the phrase “I believe,” not all intentions are expressed with the phrase “I intend.” Thus we maintain that typically outright intentions are expressed with unqualified assertion, and that consideration of expressions of intention reveals a close parallel to expressions of belief. This is straightforwardly explained by our strong cognitivism.

2.3 Intention and Planning

We turn to our third argument in defense of strong cognitivism: Intentions serve a belief-like role in planning.²³ That is, it seems to be a requirement of practical rationality

non-cognitivist view like Davidson’s, are intentions essentially ineffable? (In trying to express them, one always misses them by a little bit and ends up expressing a nearby belief!)

²⁰ For this terminology, see Holton (2008; 2009, ch.2; 2014), though we understand the notion of a partial belief somewhat differently than Holton does. For roughly the same point as we make here, see Williams (1973) and Adler (2002, Ch.7). For a different view, see Hawthorne, Rothschild and Spectre (2016).

²¹ A point made by Williams (1973, 138).

²² Thus we agree with Doyle (MS) that assertions beginning “I believe ...” and “I intend ...” are, in general, “derivative or secondary” forms of the expression of belief and intention, respectively.

²³ Holton also proposes that beliefs play an intention-like role in planning and that, indeed, intention can serve as the model of belief (2014). We think that this supports our strong cognitivism.

that a person who intends to act in a certain way should plan her conduct around the assumption that she will do this. Indeed, on Michael Bratman's influential account of intentions, with which we largely agree, the main functional role of intentions is to serve as elements in our plans (Bratman 1987). (Of course, we disagree with Bratman's rejection of strong cognitivism—a point we address in §3.5.) For example, if Petra intends to be in Paris in February, she will arrange to meet friends there, rather than in New York, and she will book a bike tour to Versailles rather than to Long Island. Such planning makes sense if Petra believes that she will be in Paris in February, but it is puzzling if she does not believe this but instead considers it an open question whether she will be there, or believes that there is a good chance that she won't.

In this respect, intentions are importantly different from desires. Desires, unlike intentions, do not rationally require any planning on our part. If Petra desires to go to Paris—and even desires this very, very much—she is in no way irrational if she fails to plan on going there. Indeed, she is not even irrational if in light of her desire she plans on *not* going there—if, say, she considers the desire to be an instance of temptation. By contrast, if Petra *intends* to go to Paris, then something is amiss if she doesn't plan accordingly. This is because intentions, unlike desires but just like beliefs, are constituted by *commitments*.²⁴ When we desire something, nothing is settled. But when we intend to do something, just as when we believe something, we have made a commitment: We have settled a question or reached a conclusion (Hieronymi 2006). It is in light of having settled this question that we make plans and settle further questions. We think it is plausible that the commitments that constitute intentions are *truth-commitments*—the truth-commitments of an agent (cf. §1.2 above).

Once again, we do not deny that the role of intentions in planning could be explained without appeal to strong cognitivism—after all, as we acknowledged earlier, our account of this role is due to Michael Bratman, who is famous for defending a form of non-cognitivism (1987; 2009a). We postpone to §3.5 a detailed discussion of Bratman's views, where we will argue that his non-cognitivist position conceals deeply cognitivist commitments. Our argument here is that the belief-like role of intentions is *especially* easy to understand if intentions are a kind of belief. Those who deny this simple view will have to appeal to something more complicated to explain why intentions play this belief-like role.

We turn now to consider whether there is any good reason to adopt such a more complicated view.

3. Objections

We argued in Section 1 that insofar as an agent's future is up to her, what she will do depends on her own beliefs about this matter, and so it is impossible for her to form

²⁴ Sinhababu (2013) offers an account of intentions as desires embedded in a network of beliefs about how one will act. We hold that this does not adequately capture the commitment-constituted character of intentions and that the extent to which it does capture it, it is the network of beliefs, rather than the desires, which explains the commitment.

these beliefs in a way that only reflects evidence about what is going to happen anyway. It is thus an essential part of being an agent that one can form practical beliefs—beliefs held in light of, and made rational by, practical reasoning. We furthermore argued that once the notion of a practical belief is recognized, it is plausible to identify practical beliefs with intentions. This becomes even clearer when we notice that, as we argued in Section 2, the characteristics of practical beliefs overlap with many of the defining characteristics of intentions: They are formed in light of practical reasoning (§2.1), they are expressed by saying what one will do (§2.2), and there is a rational requirement to treat them as “fixed points” in planning behavior (§2.3). The thesis of strong cognitivism offers an elegant and straightforward explanation of why this would be so. Given this, the question should be not what reason there is to think that practical beliefs *are* intentions, but what reason there is to think that they *are not*. Absent some such reason, we should not settle for anything weaker than our strong cognitivism.

In this section we consider several arguments philosophers have given that purport to show that intentions cannot be beliefs. The arguments are: that it is possible to intend to do things that one does not believe one will do (§3.1); that strong cognitivism leaves no way to make a principled intending/foreseeing distinction (§3.2); that beliefs are propositional attitudes, but intentions have infinitival content (§3.3); and that intention plays a motivating role in behavior that mere belief cannot (§3.4).²⁵ As we show, none of these arguments are successful. Finally, we consider whether our dispute with a non-cognitivist view like Bratman’s is merely verbal (§3.5).

3.1 *Partial Intentions*

It is often held that cognitivism is false because it is possible for intention and belief to come apart, as a person may intend to do something without being sure she will follow through. The objection is usually presented by appeal to variations on some well-known cases: For example, Michael Bratman intends to stop at the bookstore on the way home while doubting that he will, because he believes that once he gets on his bike, he will likely go into auto-pilot (1987, 37);²⁶ and Al Mele’s free-throw shooter intends to hit the basket despite believing there is a good chance that he will miss it (1992, ch.8). In reply, our contention will be that these alleged counterexamples to strong cognitivism are not only consistent with our position, but actually support it. That is because they exemplify cases in which intending, just like believing, can be *partial*.²⁷ Indeed, this is an extension of the argument concerning expressions of intention that we put forward earlier in §2.2. To develop this point fully, however, we need first to distinguish between two different forms of “partialness.”²⁸

²⁵ Another anti-cognitivist argument, that intentions cannot be beliefs since beliefs are based on evidence and evidence alone, was addressed in §1.2.

²⁶ See also Holton (2008; 2009).

²⁷ See Harman (1976, 432–3) and Setiya (2008, esp. 396) for arguments in this vein. Bratman (1987) uses “partial” to describe intentions whose details have not been filled in; this is not the sense in which we use the term here.

²⁸ Thanks to Luca Ferrero and Michael Titelbaum for helpful discussion of the argument of this section.

Let us begin with the common idea that belief can be *weak* in the sense that we can believe something without being sure that it is true—without having completely made up our mind.²⁹ (Weak beliefs are not credences; weak beliefs can be true or false, whereas credences can merely be accurate or inaccurate. Credences are not our topic.³⁰) For example, Jack might believe that the cathedral is to the right but not be sure if he remembers this correctly—since this is the sort of thing he tends to misremember. If asked by a tourist for directions, Jack might say, “I believe the cathedral is to the right, but I don’t know for sure.”

The very same thing is true of intentions as well, as we will often intend to do something without fully intending to do it—again, without having completely made up our mind. For example, Annie might intend to go to Peter’s party but be open to staying home after all. If asked by a friend whether she will go to the party, Annie might say, “I intend to go, but I am not sure I will; I may just end up watching Netflix.” This sort of weakness in Annie’s intention parallels the weakness in belief described above: Here, the reasons in favor of going to the party don’t lead her to the conclusion that she will definitely go, just as in the case above she doesn’t regard her evidence as enough to be sure that the cathedral is on the right. And just as a weak belief can be expressed with the phrase “I believe” or “I think,” a weak intention can be expressed with the phrase “I intend” or “I think”: If asked whether she will go to Peter’s party, the most Annie can honestly say is that she *intends* to go or *thinks* she will, thereby conveying that she hasn’t completely made up her mind.

The second way belief can be partial is when the lack is within the *content* of one’s belief rather than the strength of one’s attitude toward it. Suppose that Dorothy believes that Petra will fly to Paris tomorrow. While Dorothy may be entirely confident in what she believes, it is unlikely that she believes that Petra will fly to Paris *no matter what*.³¹ And this is because she will be aware that there could be events—a volcano outbreak in Iceland, say—that would prevent Petra from flying. The possibility of events like these is not usually relevant to the question of whether Petra will fly to Paris,³² so even if Dorothy accepts that they could arise, she will probably express her belief through outright assertion, assuming that her audience will share her sense of what is relevant. Still, Dorothy believes something less than if she believed that there weren’t any possible circumstances that could keep Petra from flying to Paris. Relative to the belief that Petra will fly to Paris *simpliciter*, Dorothy’s belief that Petra will fly to Paris is partial because it is *incomplete*.

We propose that intentions, too, are very often incomplete in this way, because they have a *conditional* content that is not always evident in how they are expressed. Following Luca Ferrero (2009), we distinguish two ways that the deep structure of intentions can be conditional. Consider Petra’s intention to fly to Paris. First, even if she

²⁹ “Weak belief” is the terminology of Hawthorne, Rothschild and Spectre (2016).

³⁰ Indeed, we share Holton’s (2008; 2014) reservations about credences.

³¹ It is tricky to state this point, since the domain of quantification over possibilities will typically be restricted so that he, effectively, ignores them and thus would deny that there are any such possibilities. See Lewis (1996).

³² They can become relevant, if they gain salience. We think that which possibilities are relevant, and which are properly ignored, could be explained in Lewisian fashion (Lewis 1996).

is fully committed to making the trip, Petra will know that there could turn out to be circumstances, such as a volcano outbreak along her route, that mean she won't be able to do this after all. Given that these are circumstances in which flying to Paris would be impossible, Petra's intention to fly to Paris doesn't extend to them; she doesn't intend to fly to Paris *no matter what*, but only so long as these interfering circumstances don't obtain. Possibilities like these may not be relevant in most contexts, so even if Petra accepts that they could arise she may often express her intention by asserting outright that she is going to fly to Paris. Nevertheless Petra's intention is, as Ferrero puts it, *preconditional* on certain things, as she understands that she won't be able to carry it out if the world doesn't cooperate.

The second way intentions can have a conditional structure is when their content contains what Ferrero calls *restrictive* conditions, or conditions specifying when the agent will *choose* not to act in a certain way. For example, Annie might intend to go to Peter's party *if he invites her*, or *if Marko will come along*. These are different from weakness in your intention: If the circumstances don't obtain and Annie doesn't go, it won't be because she changed her mind. (By contrast, if Annie had a weak intention to go to the party but then stayed home to watch Netflix, she will thereby have abandoned her intention.) But restrictive conditions are also different from preconditions insofar as not being invited to the party, or accompanied there by Marko, won't *prevent* Annie from going in the way a volcano outbreak would prevent Petra from flying to Paris; they are rather the circumstances in which she would voluntarily choose to do something else. While restrictive conditions like the ones so far considered are likely to be relevant to the way Annie expresses intention, other such conditions may not be: For example, it's likely that she intends to go to the party only if she doesn't come down with the flu, doesn't have a death in the family, and doesn't fall madly in love with someone who wants to go elsewhere. (Once again, these conditions are restrictive rather than preconditional: Falling in love or coming down with the flu won't *prevent* Annie from going to the party, but will lead her to *choose* not to go, by making it look like a bad idea.) Such restrictive conditions are usually taken for granted, and so they may not keep Annie from expressing her intention simply by saying that she'll go to the party. But still they mean that she intends something less than she would if her intention didn't have this conditional character, since the scope of her intention is limited by the fact that there are circumstances in which she might choose not to go.

Following Ferrero (2009, 720), the conditional "deep structure" of an intention to φ can be expressed with the following formalism:

(DS) x intends to [PRE $P_1, P_2 \dots P_m$] (φ IF $C_1, C_2 \dots C_n$)

Here, P_1, P_2 , and so on are the *preconditions* that must obtain in order for the agent to act as she intends, and C_1, C_2 , and so on are the *restrictive conditions* under which she intends to act. Thus in the example above, Annie's intention whose surface form is simply "to go to Peter's party" has the deep structure:

to [PRE the party is held, there is no large explosion nearby, etc.] (*go to Peter's party* IF I am invited, I don't have the flu, Marko will come along, etc.).

Unlike the weakness described earlier, the fact that Annie's intention has this implicitly conditional character does not mean she is likely to abandon it: She may intend *to go to the party if it takes place and she doesn't have the flu* without thereby being irresolute in this commitment. In this case it is the *content* of Annie's intention, rather than the strength of her attitude toward it, that makes her intention partial relative to another intention, also liable to be expressed by saying "I will go to Peter's party," whose content is more encompassing.

In light of this, let us return to the objection at hand: Is it possible to intend to do something that one does not believe one will do? Our reply is that all the examples that are supposed to illustrate this possibility are really only cases where a person *partially* intends to do something that she *partially* believes she will do. However, as we have seen, there are several distinct ways in which intentions can be partial, and thus there are several distinct grounds for doubting whether one will act as one intends to: (1) concerns that one will change one's mind about whether to do this, (2) recognition of the restrictive conditions that limit the scope of this intention, and (3) recognition that executing this intention will depend on preconditions that may not obtain. Let us discuss each possibility in turn.

First, a person might not believe that she will act in a way she intends to, because she recognizes that she may abandon this intention in the face of difficult or tempting circumstances. Our position is that in such a case the person only *weakly* intends to act in this way, as her attitude toward what she intends to do is not sufficiently resolute to qualify as a strong intention. To be fully resolute in her intention to do something, an agent must be committed to doing this even in the face of difficulty or temptation.³³ The person who thinks she may choose to give in to temptation has not formed an intention of this sort, precisely because she is open to reconsidering her choice when she is tempted later on. The same goes for the person who thinks she may choose to abandon her intention when the action becomes difficult or costly: Unless she thinks it might become simply *impossible* for her to act in this way, her openness to giving up in the face of difficulty reveals a limit in the strength of her intention. And this, in our view, is the same as the limit in the strength of her practical belief: A person who believes that she may change her mind about doing something only *weakly* believes that she will act in this way, and this weak belief is identical to her weak intention.

The second sort of case, in which a person's lack of a full belief that she will act in a way she intends to reflects the restrictive conditions in her intention, is similarly straightforward. If Petra intends to fly to Paris only so long as she doesn't have the flu, then her intention to fly to Paris is partial because it is incomplete: Petra doesn't intend to fly to Paris no matter what, but only to (fly to Paris IF she does not have the flu). If Petra isn't certain that she won't have the flu, then she won't be certain that she'll fly to Paris, either, but this is no objection to strong cognitivism, since if Petra's intention is strong then she will still have a strong practical belief *that she will fly to Paris if she doesn't have the flu*. Nevertheless Petra might express her state of mind by saying "I intend to fly to Paris, but am not sure if I will, since I may have the flu," because while it's natural to

³³ On the practical importance of resolutions, see Holton (2009).

assume that Petra's intention to fly to Paris is conditional on her health (so this much is not worth mentioning), that she might have the flu is not always taken for granted, and so if for some reason this possibility is salient, it ought to be mentioned explicitly. But these ways of expressing our plans in conversation with others do not show that we can have strong and complete intentions without correspondingly strong and complete practical beliefs. Indeed, someone other than Petra could hold the same combination of partial beliefs in light of the same considerations: Petra's friend Dorothy might say, "I believe Petra will fly to Paris, but I am not sure if she will since she may have the flu." And of course there is no irrationality in this complex attitude. Thus the second sort of case does not provide reason to reject strong cognitivism, but only to hold that the agent's intention is merely partial.

We turn now to the third sort of case, in which a person does not believe that she will act in a way she intends to, because of her concern that whether or not she will act in this way depends on certain preconditions that may turn out not to be satisfied. For example, suppose Petra intends to fly to Paris in February but is worried that the airline workers will be on strike, and so does not strongly believe that she'll be able to fly after all. We analyze this as a case where Petra intends to [PRE the airline is not on strike, etc.] (fly to Paris in February), and we hold that this intention is identical to her practical belief that she will fly to Paris *if the airline is not on strike*. Here, the relevant conditional element in Petra's intention is preconditional, not restrictive: Petra views the strike as something that would *prevent* her from flying, rather than a circumstance in which she'd *choose* not to go. This analysis of the content of Petra's intention is supported by the fact that if the possibility of an airline strike is sufficiently salient, then Petra won't express her intention to fly to Paris in the way that strong or outright intentions are usually expressed (see §2.2), namely by saying that this is something she *will do*. Instead, as long as Petra is mindful of the potential difficulties she'll say only that she *intends* to do this or *thinks* she will. And it is also supported by the fact that Petra won't *plan* on flying to Paris in the way we plan around things we strongly or outright intend to do (see §2.3): She won't, for example, make hotel reservations that can't be canceled in the event that she doesn't have a flight. And finally, Petra won't treat her belief about whether she'll fly to Paris as transparent to the question of what is worth doing, since she sees that whether or not she flies is not entirely up to her (see §2.1): The most she can conclude through her practical deliberation is that she will fly to Paris unless there is a strike. All of this supports the identification of Petra's incomplete intention to fly to Paris with the practical belief that she will fly to Paris as long as there is not a strike.

This preconditional analysis also applies to cases where an agent's lack of confidence in what she will do arises from a concern that she may simply fail to execute her intention.³⁴ The case of the free-throw shooter is like this: If Isaac is an accurate but imperfect shooter who is committed to hitting a shot, it will not be rational for Isaac to believe outright that he will hit it. His intention, too, will be preconditional, but in this case the relevant condition will not concern an *external* circumstance that might prevent him from making the shot, but the *internal* possibility that he might accidentally shoot in

³⁴ We are grateful to an anonymous referee for pressing us to address this sort of case.

the wrong way. This may be surprising but is not implausible: We understand our control over ourselves—and especially our bodies—to be limited. Because the limitations are internal, we don't think of the resulting failures as a matter of being prevented from acting as we mean to. But we also don't think of such failures as entirely inexplicable. Isaac can miss a free throw, not because it is blocked or blown by the wind, but *because of how he shoots it*—and whether or not he shoots it in a way that misses is not something that Isaac has total control over. He can, then, rationally intend to hit a free throw only in a way that is preconditional on his not shooting it wrongly.

In another set of cases, a similar possibility of failure arises due to limits in our psychological self-control rather than our control over the movements of our bodies. Sometimes we forget things—we go on “auto-pilot,” like Bratman who intends to stop at the bookstore but forgets (1987, 37)—or we blunder, or we lose focus in ways that keep us from doing what we intend. In such cases, our failure to do what we intend is explained by lapses in the “executive control” by which we keep these kinds of things from happening. And it is plausible to see our intentions as preconditional on our exercising this control successfully—a condition which we cannot expect to have total control over. But there is no reason to see the preconditonality of intentions on the absence of accidental executive failures as categorically different from the preconditonality of intentions on the absence of external conditions that would prevent our success, such as volcanoes or strikes or simple bad luck. Though the question may arise whether a person who believes that she will fail in these ways can take further steps to guard against this possibility, and, if so, how we should regard her intention if she does not take them, in itself the worry that an involuntary lapse may prevent one from executing an intention does not show that the intention itself is weakly held.

We conclude that consideration of partial intentions does not succeed in undermining strong cognitivism. Rather, it supports it by showing how the strength and extent of our intentions and beliefs are rationally required to fit together into one consistent worldview.

3.2 *Intention and Foresight*

We turn now to the second objection. Michael Bratman has argued that strong cognitivism cannot capture the difference between what one intends and what one foresees as a side-effect of one's intentional action (Bratman 1987; 2009a; and cf. Paul 2011). For example, suppose Petra intends to fly to Paris and foresees that doing so will contribute to global warming. Intuitively, Petra can believe that she will contribute to global warming without intending to do this. According to strong cognitivism, however, Petra's intention to fly to Paris is identical to her belief that she will do so. What, then, is the difference between this intention and the belief that is her foreseeing that she will pollute? How can we maintain the distinction between intention and foresight if both are just a matter of belief?

Our reply is that according to our view one intends to do something if and only if one believes in light of one's *practical reasoning* that one will do it. By contrast, one foresees, or predicts, that one will do something only if one believes in light of *evidence*

that one will bring it about, that one will do so. Thus in the case at issue, Petra intends to fly to Paris, and hence believes that she will fly to Paris, in light of reasoning that shows this to be worthwhile. She also believes that by flying to Paris, she will contribute to global warming, because she has evidence that flying contributes to global warming. The difference in the *grounds* of these beliefs is what makes one of them an intention, or practical belief, and the other not. Hence, as long as we keep sight of the different forms of reasoning that ground belief, we will be able to draw the intending/foreseeing distinction appropriately.

This might seem insufficient.³⁵ It may seem impossible to isolate when a conclusion is grounded in practical reasoning and when it is grounded in theoretical reasoning, since in making a decision we have to consider not just what we have reason to do but also what we can expect to happen if we go ahead. For example, in deciding whether to take the flight, Petra should consider not just whether she wants to go to Paris but also whether this justifies the foreseeable consequences of doing so—for instance, whether the air pollution is a cost worth incurring. If she had reason to expect that her taking the flight would have sufficiently terrible consequences—e.g., that it would bring about environmental disaster all by itself—then she would not see it as worthwhile to take the flight. Thus it may seem that in choosing to take the flight, Petra also chooses to bring about the expected consequences. And this can make it appear that Petra’s belief that she will contribute to global warming, no less than the belief that she will go to Paris, is grounded in her practical reasoning.³⁶

We hold that the answer to this objection turns on the instrumental or calculative nature of practical reasoning, in virtue of which an agent’s practical reasoning settles the question of *how* she will achieve her ends (see §1.1). This does not mean that to form an intention to do something, an agent has to formulate a plan in all its particular details.³⁷ Rather, the point is that insofar as one reasons practically, the question “How will you do that?”, no less than Anscombe’s “Why?”-question, finds application. An agent’s practical beliefs thus comprise a causal order, and the form of causality is that of instrumental teleology: One who believes that she is going to ϕ (drive to the theater) because she is going to ψ (see a movie) understands her ϕ -ing not as something that will be *produced* by her ψ -ing, but rather as a suitable *means* to it. And we identify this instrumental order in the agent’s view of her future with the content of her intentions: It is in light of regarding ϕ -ing as an appropriate means to ψ -ing, and ψ -ing as a worthwhile end, that a person has a practical belief that she is going to ϕ because she is going to ψ . Any conclusion of practical reasoning will be such as to have a place within such an order of means and ends. And the content of an agent’s view of her future will encompass this order within itself: One intends to ϕ because she is going to ψ ; to accomplish this by doing α , β , and γ ; to try doing δ instead if certain likely difficulties arise; and so on.

³⁵ Thanks to Michael Bratman and Sarah Paul for pressing us on this issue.

³⁶ This issue has been at the core of discussions about the relation between intention and beliefs. See Harman (1976) and Bratman (1987) for the *loci classici* and Rozeboom (2015) for recent discussion.

³⁷ A point stressed by Bratman (1987); and cf. Harman (1976, 446–9). But cf. also Ford (2016) on the “particularizing” function of practical thought.

The instrumental structure of practical thought is what allows our strong cognitivism to resist this renewed objection. We propose that what distinguishes someone's intended actions from the foreseen side-effects of her actions is the way they relate to this order of means and ends: The thought that, say, one will buy a ticket because one is flying to Paris involves a different kind of causal rationalization than the thought that, because one is flying to Paris, one is going to contribute to global warming. The "because" of practical reasoning embeds the belief that an agent will do something in an order of instrumental rationality which rationalizes this belief as a representation of a suitable means to an end. By contrast, the "because" of theoretical reasoning embeds a belief in an order of theoretical rationality and thus a structure of productive or "efficient" causation which rationalizes this belief in the manner of a prediction. Even if we reach practical and theoretical conclusions at once, in one act of the mind and in mutual interdependence, they remain distinctly rationalized. And since we don't propose to collapse practical and theoretical reasoning, we can, in principle, distinguish their conclusions in this way.

The difficulty in question arises because it seems that in deciding to do something we must, as it were, reason in one act of the mind in a way that is at once practical and theoretical: We must both determine what we will do and also predict what will (likely) happen as a result of our doing it. And this is exactly as it should be, since in settling what we will do, we also settle what thereby happens. This means that it would be irrational to ignore in our practical deliberation what will (likely) happen as a consequence of our possible actions. However, it does not show that our beliefs that these things will happen are grounded in practical reasoning. Rather, an agent's conception of the future is rich enough to distinguish those elements in it that are embedded in the teleological order of means and ends from those that lie outside this order, and are seen by the agent as *effects* of her chosen actions rather than as means to performing them.

3.3 *Intention as a Propositional Attitude*

The third objection we consider is that belief is a propositional attitude but intention is not, as an intention does not have a propositional content.³⁸ Moreover, because beliefs are propositional attitudes they can be false—like suppositions and stipulations. But because intentions seemingly cannot be false, it seems that they are not propositional attitudes and therefore cannot be identical to beliefs.

In response we hold that intentions are indeed propositional attitudes. This is plausible because, as we argued earlier (§2.1), outright intentions are expressed through outright assertion. For example, if you intend to go to Peter's party, you will express your intention by saying, "I will go to Peter's party." In so doing, you express an attitude whose content is a proposition—and what you say will be *false* if you do not do what you said you would.³⁹ According to strong cognitivism, the proposition that is the content of

³⁸ See especially Thompson (2008) and, more recently, Hornsby (2016) for this argument. For an opposing argument with which we have some sympathy, see Ferrero (2013).

³⁹ It can be tempting to read Anscombe as denying this, but she does not: She writes in *Intention* that the expression of an intention might be such that "what [one] says is not true," and that in such a case "there is a discrepancy between the language and that of which the language is

what you say in expressing your intention in this way is also the content of your intention itself.

Though it is hard to deny that verbal expressions of intention are truth-evaluable in the way we have just described, we concede that it is awkward to speak of *intentions* as true or false. And this puts pressure on our claim that they are propositional attitudes. However, we propose that the unnaturalness of *saying* that an intention is false can be explained by recalling the distinction between practical and theoretical beliefs. The claim that a certain statement or attitude is false will usually⁴⁰ be based on evidence. But as we have argued (§1.1), from the agent's perspective evidence about the likelihood of her performing certain actions does not settle the question of what she will do: A person will often take evidence that she may not φ not as a reason to abandon her intention, but rather as a reason to find means of overcoming the difficulty at issue. Because of this, if one wishes to contradict or "falsify" an intention in a way the agent will have to regard as relevant, the way to do this is by questioning whether the intended action is good or worthwhile. This is an attempt to persuade the agent that her intention is untrue insofar as it is an attempt to persuade her, through practical reasoning, to conclude that she *will not do* what she presently believes she will do, though the distinctively non-evidential nature of this challenge makes it natural not to construe it as a matter of truth and falsehood.

The other source of pressure against the view that intentions are propositional attitudes comes from the observation that the verb "intend" usually takes an infinitival construction as its object. (Not always: Both authors intend that they finish work on this paper, and that their children learn to read.⁴¹) But one should be careful before drawing conclusions about the metaphysics of mind simply from the surface grammar of English. Moreover, we think that this usage can be seen as well motivated even if intentions have propositional content. A first reason to use "intend" in this way is that, as we emphasized above (§1.2), practical reasoning is essentially the reasoning of an agent who is thinking about the future insofar as it is up to her, and it is reasonable to conclude that something will be true in light of its appearing as worthwhile or good only insofar as the truth of this conclusion turns on what one thinks about it. This means that the possible objects of practical deliberation are essentially tied to one's own actions. And that is why it can go without saying that when I intend to go to Peter's party, the content of my intention is that *I* will go to it.⁴²

A second motivation for this usage is that in having the attitude that in English is called *intending to do* something, a person does not relate to her intended action as an

a description" (1957/2000, 56, 57). In her view cases like these are such that the *fault* or *mistake* lies in action rather than judgment, but this does not mean the judgment is not false. For discussion of these difficult passages, see Schwenkler (2015, 20–2).

⁴⁰ Usually, though not always: Sometimes, as Anscombe notes (1957/2000, 55), a statement of the form "No, you will not φ " may *itself* be the expression of an intention—that is, the intention to stop this person from acting in such a way. In such a case the statement is what Anscombe calls the "head-on contradiction" of the corresponding intention (1957/2000, 92).

⁴¹ As Hornsby (2016, 5) points out, what follows "that" in sentences like these is not a propositional clause, since the verb is subjunctive.

⁴² Thus we accept Boyle and Lavin's (2010) argument that any concept of φ -ing that could be part of the content of an intention will presuppose a *de se* presentation of the agent (cf. Schroeder 2012, 463). See Ferrero (2013, 77) for an opposing view.

event that she intends to bring about.⁴³ For example: Suppose Niko does not intend to study tomorrow; he knows that he should, but simply doesn't feel up to the task. Knowing this, Niko might decide to take a pill that will make it the case that, tomorrow, he will feel differently and so form an intention to study. A possible way to describe this case is as one in which Niko intends *that he study tomorrow*, but does not (yet) intend *to study tomorrow*. One way to account for this difference would be if the attitude of *intending to do* something has a distinctively non-propositional content. But another strategy, which we favor, holds that the difference between these two attitudes lies in whether or not the agent intends to do something in order to *bring it about* that one does something further. In the case just described, Niko intends to take a pill, an intended consequence of which will be that he studies intentionally later on. This is different from the case where Niko simply intends, now, to study (either now or later). But the difference is not that one representation is propositional and the other not.⁴⁴ Rather, the difference lies *within* the propositional contents of the two representations: In the first case, Niko intends (and so believes) *that he will take the pill, and thereby bring it about that he will study*; and in the second, Niko simply intends (and so believes) *that he will study*. The content of a simple intention to φ does not contain an intention to act upon oneself in order to produce that behavior. We propose that this simplicity is marked in English with the use of the infinitival construction.

3.4 Intention and Motivation

The final objection we consider is that intention cannot be belief because intending to do something involves being *motivated* to do it in a way that merely believing one will do something does not. A quick version of this objection betrays a misunderstanding of our position: It says that if a person already believes she is going to do something, then this confidence will obviate the need for formulating a plan of action and working to overcome barriers to success. The misunderstanding behind this objection is that, as we emphasized above (see §3.1), in intending to do something a person will not believe *simpliciter* that she is going to do it—for this belief may be conditional on the belief that she is going to do certain other things that she sees as a useful or necessary means to this end. That is, a person may strongly intend to do something while understanding that she will be able to carry out this intention only under certain conditions—in particular, only if she follows a plan and is prepared to surmount obstacles as they arise.

A more difficult version of this objection is that the *form* of motivation involved in the intention to do something shows that intention must be different from, or at least more than, a merely cognitive state like belief. Consider the case of Anna Karenina as

⁴³ See Boyle (forthcoming). We are grateful to Matthew Boyle for illuminating discussion of what follows.

⁴⁴ In light of this, while we agree with Jennifer Hornsby's claim that a sentence of the form "She intends that she φ " appears to introduce a sort of distance between intention and its fulfilment—a distance never heard when "She intends to φ " is said' (2016, 5), we do not accept her conclusion that this shows the content of "intending to" cannot be given a propositional analysis.

discussed by Sarah Paul:

Take Anna, who upon reflecting on all the reasons why her love affair with V is a terrible thing for her reputation, her children and so forth decides to break it off with him. Someone with thorough knowledge of her motivational profile—her passion for V, her loathing of her husband, her feelings of suffocation in the role of political housewife—would conclude that her motivation to end the affair is far too weak to bring her even to try to break it off. Knowing this we may be tempted to say that Anna never really intended to break up with him. She believes she ought to and has even made the decision to but was at no point committed to the breakup in the right way to count as intending it. (Paul 2012, 338)

Tailored to the present debate, Paul’s case might be seen as one where Anna has a practical belief that she is going to break it off with V (perhaps this is what Paul calls her “decision”), but lacks the intention to do so. On this analysis, what a person intends is related to her emotional and desiderative profile, and not merely what she believes that she will do.⁴⁵

We deny, however, that this is the only possible description of this case, or even the most natural one. This is because the very same reasons in favor of saying that Anna does not *intend* to break it off with V are also reasons for saying that she does not *believe* that she will do this, either. If Anna is asked whether she is going to break it off with V, then she *may* say that she is going to, and may take it that this is what she really believes—but there will be circumstances (e.g. if she is asked the question by V himself) where she probably will not give this answer, and her broader planning behavior will show that she does not treat this decision as a fixed point in her view of the future. More generally, it is not only in regard to desires and intentions that a person’s desires and emotions are relevant to determining what she thinks: Especially when matters of human importance are at stake there will be more to believing something than just having a representation in a certain cognitive “box,” and to the extent that the question of whether one believes something bears significantly on one’s self-conception (as it surely does in Anna’s case) it will be easy to fall into error in what one thinks about her own beliefs.⁴⁶

The idea that belief is purely intellectual and dispassionate, that the criteria for belief cannot require commitment and motivation, and that what a person believes floats free of how things are with her “deep self,” assumes a view of belief that has already been discredited in other domains. Our arguments for the existence of practical beliefs provide yet another reason to reject this view.

3.5 *A Verbal Dispute?*

We will now consider one final concern, namely that our cognitivist position may

⁴⁵ We are grateful to Philip Clark for raising this objection and suggesting that the competing view be stated in this way.

⁴⁶ For important discussion of these last two points, see McGeer (2007) and Schwitzgebel (2010).

seem to be merely a verbal variant of a non-cognitivist view like Bratman's. Indeed, we draw on Bratman's central thesis that the main functional role of intentions is to serve as elements in our plans (1987). And we maintain that intentions are distinct from theoretical beliefs—since we hold that they are beliefs that are based on practical reasons. Couldn't the non-cognitivist concede this point but continue to reserve the word "belief" for what we call theoretical beliefs and hold that intentions are different in kind from beliefs?⁴⁷

In reply, we accept that our view could be regarded as a verbal variant of a non-cognitivist view like Bratman's—but only because Bratman's alleged non-cognitivism is in fact deeply cognitivist.⁴⁸ Below we explain why this is.

To make our case, we first need to distinguish two different senses of "non-cognitivism." In one sense, non-cognitivism is the view that the consistency and coherence requirements on intentions—such as the instrumental principle, according to which it is a rational failure to fail to take what one believes to be the necessary means to an end—are not derivable from the consistency and coherence requirements on beliefs. We have no objections to this kind of non-cognitivism. Indeed, since we don't propose to collapse the distinction between practical and theoretical reasoning, we actually endorse it. And we suggest that this is the kind of non-cognitivism that should be of central importance to Bratman.

In another sense, non-cognitivism is the view that intentions don't consist in or entail beliefs. It is this kind of non-cognitivism that we reject. Moreover, we hold that this kind of non-cognitivism cannot do justice to the planning role of intentions (contrary to Bratman's view!). Indeed, we think that Bratman holds an implicit cognitivist commitment, as we will now argue.

To account for the role of intentions in plans, Bratman argues that it is a rational requirement that one's intentions be consistent with one's beliefs. He writes,

Central to the planning theory is the idea that intentions—in contrast with ordinary desires—are both embedded in characteristic regularities and are subject to distinctive rational pressures for consistency and coherence. There is, in particular, a rational demand that one's intentions, taken together with one's beliefs, fit together into a consistent model of one's future (2009a, 29).

Much of Bratman's recent work is concerned to explain the normative force of such a rational requirement (2009a; 2009b). Bratman seeks to defend it as a wide-scope rational requirement that derives its normative force from the practical advantages of self-governance, since "consistency and coherence of intention is a necessary constitutive element in the ... self-governance of planning agents like us" (2009b, 431).

Without taking issue with Bratman's larger view about the normative force of

⁴⁷ Doyle (MS) quotes Wittgenstein (1986, §79) in this connection: "Say what you choose, as long as it does not prevent you from seeing the facts."

⁴⁸ Another officially non-cognitivist view that strikes us as deeply cognitivist is Holton's (2008; 2009, ch.2; 2014)—especially in light of Holton's arguments that intentions can serve as a model for belief, because they share their central functional roles.

rational requirements, we hold that Bratman's claim about the rational requirement of *consistency* carries with it a cognitivist commitment. That is, the claim that it is a *rational* demand that intentions and beliefs fit together into a consistent model of our future implicitly implies that intentions and beliefs are governed by the same standard of correctness. And if the standard of correctness is that which distinguishes a mental state as a belief, this implies that intentions must be beliefs.⁴⁹

To see this, consider a mental state that does not yield irrationality if it is inconsistent with one's beliefs—for instance an act of imagination. We can imagine that it is raining while we believe that it is sunny, without any irrationality. Why? Because imagination is not subject to truth as a standard of correctness: An act of imagining is not incorrect if what is imagined is false. Imagination is not subject to the same standard of correctness as belief; that is why inconsistency of imagination with belief is not a rational failure. The same is true of desire and other non-cognitive attitudes: Their inconsistency with beliefs does not yield irrationality, because they do not have truth as their standard of correctness.

If, however, there is a rational demand that intention and belief be consistent, then the two states have to have this much in common: They are subject to the same standard of correctness. Now, Bratman does not deny that belief is subject to truth as its standard of correctness; he is not a pragmatist in the ethics of belief. That is why his claim that there is a rational demand that intentions and beliefs be consistent carries an implicit commitment to *truth* being the standard of correctness of intentions as well—a commitment that is easily explained by the thesis that intention is practical belief.

This simple explanation of why intention and belief are subject to the same standard of correctness could be challenged, of course, by showing that it is not always irrational to intend without believing. Bratman argues that such *intention-belief incompleteness* can come in two forms (1987, 37–9): A person may either (i) intend to do something without having considered the question of whether she will do it, and thus lack any belief at all about the matter, or (ii) intend to do something while having only a weak or incomplete belief that she will do it. And we have already explained why we reject both of these alleged possibilities: The first is challenged by our arguments that to form an intention *is* to settle the question of what one will do (§§2.1 and 2.3), and the second is addressed in our account of partial intention (§3.1). In light of this, we suggest that Bratman's arguments for intention-belief consistency ultimately support our strong cognitivism: What best explains the rational requirement that intention and belief fit together into a consistent model of the future is, simply, that intending is believing.

It may be surprising that we would accept, in the context of explicating our strong cognitivism, that there is such a deep similarity between our view and a paradigmatically non-cognitivist view like Bratman's. However, we believe that this is only because Bratman's alleged non-cognitivism is in fact deeply cognitivist. If “there is a rational demand that one's intentions, taken together with one's beliefs, fit together into a

⁴⁹ To clarify our dialectic with Bratman: We think that if Bratman's arguments about self-governance vindicate his claim about rational requirements, they do so because they show that self-governance requires that intentions and beliefs be held to the same standard of correctness. For further discussion of Bratman's view, in particular, see Marušić (2013; 2015, ch. 3.1).

consistent model of one's future," then intentions and beliefs are in the business of "modelling"—and so presumably representing—the same thing, namely how the world will be. Thus we ask: In what sense is this view non-cognitivist? Isn't a model or representation of how the world will be a form of cognition—especially if it is governed by truth as the standard of correctness? Indeed, if, according to non-cognitivism it is possible to form one's view of the future in light of practical reasoning rather than evidence, and to assert, plan and act in light of this view, then intentions have the same functional roles as belief. All this makes it plausible that if our dispute with a non-cognitivist like Bratman is merely verbal, this is not because our view is irresolutely cognitivist but because the other view is irresolutely non-cognitivist.⁵⁰ This, we think, takes the sting out of the objection that the dispute with the non-cognitivist is merely verbal.

4. Conclusion

It is a persistent philosophical question how to reconcile practical reasoning and theoretical reasoning. And it is a persistent philosophical tendency to collapse or to insulate the two—for instance, to deny that practical reasoning is concerned with factual matters or construe it as a special kind of reasoning in light of evidence. But this gives rise to persistent problems—especially because intentions are commitments to change factual matters and because they are importantly related to assertions and planning about what will happen. Our strong cognitivism explains the relation between theoretical reasoning and practical reasoning without insulating them, but also without collapsing either form of reasoning into the other. Both theoretical reasoning and practical reasoning are concerned with factual matters. In theoretical reasoning, we reflect on factual matters as observers; in practical reasoning, we reflect on factual matters as agents. And as agents, who are in a position to make something true, we enjoy a great privilege: We get to settle what will be true by appeal to considerations that show its being true worthwhile. We get to believe, because we see something as good.⁵¹

⁵⁰ Sarah Paul's (2009) view is more resolutely non-cognitivist, since Paul, to our knowledge, does not endorse Bratman's claim that rationally requires that intentions and beliefs be consistent.

⁵¹ We are grateful to Matthew Boyle, Michael Bratman, Phil Clark, Jimmy Doyle, Luca Ferrero, Jennifer Frey, Pamela Hieronymi, Eli Hirsch, Paul Hurley, Jon Hurst, Christian Kietzmann, Tatyana Kostochka, Michael Kremer, Clayton Littlejohn, Jennifer Marušić, Al Mele, Richard Moran, Sarah Paul, Jonathan Payton, Juan Piñeros, Piers Rawling, Sebastian Rödl, Justin Snedgar, Michael Thompson, Michael Titelbaum, Stephen White and several anonymous reviewers for helpful comments. We would also like to thank audiences at the 2015 Meeting of the Florida Philosophical Association, the 2015 St. Louis Annual Conference on Reasons and Rationality and the 2016 Central Division Meeting of the American Philosophical Association, and especially Ben Schwan and Samuel Asarnow who commented on our paper at SLACRR and the APA. Berislav Marušić gratefully acknowledge a fellowship from the Humboldt Foundation and a Theodore and Jane Norman Grant from Brandeis University. We dedicate this paper to our children: Jack, Daniel, Marko, Petra, Peter, Annie, Dorothy, Isaac, and Niko.

References

- Adler, Jonathan. (2002). *Belief's Own Ethics*. Cambridge, MA: The MIT Press.
- Anscombe, G.E.M. (1957/2000). *Intention*. Cambridge, MA: Harvard University Press, 2000.
- Bach, Kent and Robert M. Harnish. (1979). *Linguistic Communication and Speech Acts*. Cambridge, MA: The MIT Press.
- Bok, Hilary. (1998). *Freedom and Responsibility*. Princeton, NJ: Princeton University Press.
- Boyle, Matthew. (Forthcoming). "Knowledge as Subject and Agential Knowledge." In *Beings and Doings*, edited by Alexander Geddes and Alec Hinshelwood. Oxford: Oxford University Press.
- Boyle, Matthew and Douglas Lavin. (2010). "Goodness and Desire." In *Desire, Practical Reason, and the Good*, edited by Sergio Tenenbaum. Oxford University Press.
- Bratman, Michael. (2009a). "Intention, Belief, Practical, Theoretical." In *Spheres of Reason*, edited by Simon Robertson. Oxford: Oxford University Press.
- Bratman, Michael. (2009b). "Intention, Practical Rationality, and Self-Governance" *Ethics* 119, 411–43.
- Bratman, Michael. (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Broome, John. (2002). "Practical Reasoning." In *Reason and Nature: Essays in the Theory of Rationality*, edited by Jose Bermúdez and Alan Millar. Oxford: Clarendon Press, 2002.
- Davidson, Donald. (1971). "Agency." In *Essays on Actions and Events*. Oxford: Oxford University Press, 1980.
- Davidson, Donald. (1978). "Intending." In *Essays on Actions and Events*. Oxford: Oxford University Press, 1980.
- Doyle, James. (MS). "Intention as Belief."
- Falvey, Kevin. (2000). "Knowledge in Intention." *Philosophical Studies* 99, 21–44.
- Ferrero, Luca. (2013). "Can I Only Intend My Own Actions?" In *Oxford Studies in Agency and Responsibility* 1, edited by David Shoemaker. Oxford: Oxford University Press.
- Ferrero, Luca. (2009). "Conditional Intentions." *Noûs* 43, 700–41.
- Ford, Anton. (2016). "On What Is in Front of Your Nose." *Philosophical Topics* 44, 141–61.
- Ford, Anton. (2017). "The Representation of Action." In *Philosophy of Action: Royal Institute of Philosophy Supplement* 80, edited by Anthony O'Hear, 217–33.
- Grice, H. Paul. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Grice, H. Paul. (1971). "Intention and Uncertainty." *Proceedings of the British Academy* 57, 263–79.
- Harman, Gilbert. (1986). "Willing and Intending." In *Philosophical Grounds of Rationality: Intentions, Categories, Ends*, edited by Richard Grandy and Richard

- Warner. Oxford: Oxford University Press.
- Harman, Gilbert. (1976). "Practical Reasoning." *Review of Metaphysics* 29, 431–63.
- Hawthorne, John, Daniel Rothschild and Levi Spectre. (2016). "Belief is Weak." *Philosophical Studies* 173, 1393–404.
- Hieronymi, Pamela. (2009). "The Will as Reason." *Philosophical Perspectives* 23, 201–20.
- Hieronymi, Pamela. (2006). "Controlling Attitudes." *Pacific Philosophical Quarterly* 87, 45–74.
- Holton, Richard. (2014). "Intention as a Model for Belief." In *Rational and Social Agency: The Philosophy of Michael Bratman*, edited by Manuel Vargas and Gideon Yaffe. Oxford: Oxford University Press.
- Holton, Richard. (2009). *Willing, Wanting, Waiting*. Oxford: Oxford University Press.
- Holton, Richard. (2008). "Partial Belief, Partial Intention." *Mind* 117, 27–58.
- Hornsby, Jennifer. (2016). "Intending, Knowing How, Infititives." *Canadian Journal of Philosophy* 46, 1–17.
- Ismael, Jenann. (2012). "Decision and the Open Future." In *The Future of the Philosophy of Time*, edited by Adrian Bardon. New York: Routledge.
- Korsgaard, Christine M. (1986). "Skepticism About Practical Reason." *Journal of Philosophy* 83, 5–25.
- Langton, Rae. (2004). "Intention as Faith". *Royal Institute of Philosophy Supplement* 55, 243–58.
- Lewis, David. (1996). "Elusive Knowledge." *Australasian Journal of Philosophy* 74, 549–67.
- Marušić, Berislav. (2015). *Evidence and Agency: Norms of Belief for Promising and Resolving*. Oxford, UK: Oxford University Press.
- Marušić, Berislav. (2013). "Promising against the Evidence." *Ethics* 123, 292–317.
- Marušić, Berislav. (2012). "Belief and Difficult Action." *Philosophers' Imprint* 12, 1–30.
- McGeer, Victoria. (2007). "The Moral Development of First-Person Authority." *European Journal of Philosophy* 16, 1, 81–108.
- Mele, Alfred. (1992). *Springs of Action*. New York: Oxford University Press.
- Moran, Richard. (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton: Princeton University Press.
- Paul, Sarah (2012). "How We Know What We Intend." *Philosophical Studies* 161, 2, 327–46.
- Paul, Sarah. (2011). "Deviant Formal Causation." *Journal of Ethics and Social Philosophy* 5, 3, 1–23.
- Paul, Sarah. (2009). "How We Know What We're Doing." *Philosophers' Imprint* 9, 1–25.
- Raz, Joseph. (2011). *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Rozeboom, Grant J. (2016). "Side Effects and the Structure of Deliberation."

Journal of Ethics and Social Philosophy 9, 2.

Schroeder, Mark. (2012). "The Ubiquity of State-Given Reasons." *Ethics*, 122, 457–88.

Schwenkler, John. (2015). "Understanding 'Practical Knowledge'." *Philosophers' Imprint*, 15, 1–31.

Schwitzgebel, Eric. (2010). "Acting Contrary to Our Professed Beliefs, or the Gulf Between Occurrent Judgment and Dispositional Belief", *Pacific Philosophical Quarterly* 91, 531–53.

Setiya, Kieran. (2009). "Practical Knowledge Revisited." *Ethics* 120, 128–37.

Setiya, Kieran. (2008). "Practical Knowledge." *Ethics* 118, 388–409.

Setiya, Kieran. (2007). *Reasons without Rationalism*. Princeton: Princeton University Press.

Setiya, Kieran. (2003). "Explaining Action." *The Philosophical Review* 112, 339–93.

Shah, Nishi. (2006). "A New Argument for Evidentialism." *The Philosophical Quarterly* 56, 481–98.

Sinhababu, Neil. (2013). "The Desire-Belief Account of Intention Explains Everything." *Noûs* 47, 680–96.

Thompson, Michael. (2008). *Life and Action: Elementary Structures of Practice and Practical Thought*. Cambridge, MA: Harvard University Press.

Velleman, David and Nishi Shah. (2005). "Doxastic Deliberation." *The Philosophical Review* 114, 497–534.

Velleman, David. (1996). "The Possibility of Practical Reason." In *The Possibility of Practical Reason*. Oxford: Oxford University Press, 2000.

Velleman, David. (1989/2007). *Practical Reflection*. CSLI Publications.

Velleman, David. (1985). "Practical Reflection." *The Philosophical Review* 94, 33–61.

Vogler, Candace. (2002). *Reasonably Vicious*. Cambridge: Harvard University Press.

Wedgwood, Ralph. (2002). "The Aim of Belief." *Philosophical Perspectives* 16, 267–97.

Williams, Bernard. (2002). *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press.

Williams, Bernard. (1973). "Deciding to Believe." In: *Problems of the Self: Philosophical Papers 1956–1972*. Cambridge, UK: Cambridge University Press.

Williamson, Timothy. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Wilson, George. (2000). "Proximate Practical Foresight." *Philosophical Studies* 99, 3–19.

Wittgenstein, Ludwig. (1986). *Philosophical Investigations*, 3rd edition. Trans. G.E.M. Anscombe.